

N82-15027

NASA Technical Memorandum 73296

Automated Pilot Advisory System

J. L. Parks, Jr.

and

J. G. Haidt



November 1981

NASA

National Aeronautics and
Space Administration

Wallops Flight Center

Wallops Island, Virginia 23337
AC 804 824-3411

Automated Pilot Advisory System

J. L. Parks, Jr.

**NASA Wallops Flight Center
Wallops Island, Virginia 23337**

and

J. G. Haidt

**Research Triangle Institute
Research Triangle Park, North Carolina 27709**



**National Aeronautics and
Space Administration**

Wallops Flight Center
Wallops Island, Virginia 23337
AC 804 824-3411

CONTENTS

	Page
SUMMARY.	i
1.0 INTRODUCTION.	2
1.1 Concept.	2
1.2 Outline.	5
2.0 DESCRIPTION OF EXPERIMENTAL APAS.	11
2.1 Radar, Antenna, and Clutter Suppression Subsystem.	11
2.1.1 Digitizer	17
2.1.2 Minicomputer and Peripherals.	18
2.1.3 Microcomputer and Peripherals	19
2.2 System Operation.	24
2.2.1 Airport Advisories.	25
2.2.2 Traffic Advisories.	29
3.0 STRUCTURE OF EXPERIMENTAL APAS.	42
3.1 Tracking Data Unit	43
3.1.1 Radar Module.	43
3.1.2 Target Detection Module	51
3.1.3 Track-While-Scan Module	61
3.2 Weather Data Unit.	90
3.2.1 Overview.	92
3.2.2 Data Acquisition Module	93
3.2.3 Data Processing Module.	95
3.2.4 Runway Selection Module	101
3.3 Voice Response Unit.	106
3.3.1 Overview.	106
3.3.2 Speech Processing Module.	109
3.3.3 Message Formatting Module	120
3.3.4 Message Entry Module.	136
3.3.5 VRU Timing Module	136
4.0 TEST AND EVALUATION OF THE EXPERIMENTAL APAS.	140
4.1 Overview	140
4.1.1 Clutter	140
4.1.2 Timing.	141
4.1.3 Coverage.	141

4.1.4 Pilot Acceptance.	142
4.2 Low Density Traffic Test	142
4.2.1 Data Rate Tests	142
4.2.2 Radar System Tests.	143
4.3 High Density Traffic Tests	146
4.4 Synopsis of Test Results	146
4.5 Pilot Evaluation	151
5.0 PROPOSED DEVELOPMENT.	156
6.0 CONCLUSION.	157
APPENDIX	158
REFERENCES	179

AUTOMATED PILOT ADVISORY SYSTEM
John L. Parks, Jr., James G. Haidt*
Wallops Flight Center

SUMMARY

The APAS was conceived as a low cost automated system to provide general aviation pilots with air traffic and local weather information. The APAS was designed to be a natural extension of the procedural visual flight rules (VFR) system used at uncontrolled airports and, as an advisory system, was designed to enhance the see-and-be-seen rule of flight.

The experimental system described in this report was developed, tested, and demonstrated over a two-year period. Demonstration testing occurred at Manassas Municipal Airpark, Manassas, Virginia, over a seven week period. One hundred pilots who participated in the demonstration reported that they favored the APAS over a self-announcement system by better than 5 to 1. The Manassas testing demonstrated the feasibility of the APAS concept, and it indicated several areas where enhancements could be employed.

This report documents the results of the development, testing, and demonstration processes, describes the system which evolved, and recommends system enhancements to improve performance.

*Research Triangle Institute, Research Triangle Park, North Carolina.

1.0 INTRODUCTION

The high-density, uncontrolled airport presents a source of numerous operating problems and hazards to the general-aviation pilot. This fact is evidenced by the large percentage of mid-air collisions which occur in this environment, even in clear weather. Furthermore, pilots approaching such airports frequently do so lacking information concerning weather conditions, the active runway, or local traffic conditions. The anticipated growth of the general aviation fleet in the next decade will further compound these problems.

The system currently used to address this issue requires pilots to self-announce (traffic advisory) over a UNICOM radio channel and request the active runway (airport advisory) from the fixed base operator (FBO). The UNICOM radio channel is also used for general information and requests and can be shared by several different airports. For example, the UNICOM at Manassas airport is shared by Manassas, Montgomery County, Warrenton, and Freeport. The problems with this type of system are as follows:

- (1) not all pilots self-announce;
- (2) active runway information may not be available (FBO may be absent from the radio performing other jobs, etc.);
- (3) there may be radio interference due to multiple transmissions; and
- (4) self-announcement at one airport may be overheard and misinterpreted by pilots at another airport.

With state-of-the-art technology, it appears feasible to remedy this situation with a low-cost, automatic, ground-based system which would broadcast (1) traffic advisories to alert pilots to the presence of other aircraft, and (2) airport advisories to inform them of the active runway and current airport weather conditions. To this end, the National Aeronautics and Space Administration (NASA), in cooperation with the Federal Aviation Administration (FAA), has developed an experimental Automated Pilot Advisory System (APAS) to provide the indicated advisories, a system designed to enhance the see-and-be-seen rule used at uncontrolled airports. The purpose of this report is to describe APAS - its configuration, operation, and performance.

1.1 Concept

The APAS concept, illustrated in Figure 1-1, involves a computer-based system which continuously and automatically monitors air traffic and weather conditions at an uncontrolled general aviation airport and assembles the information obtained into speech for area-wide broadcasting. Successful realization of this concept implies

AUTOMATED PILOT ADVISORY SYSTEM

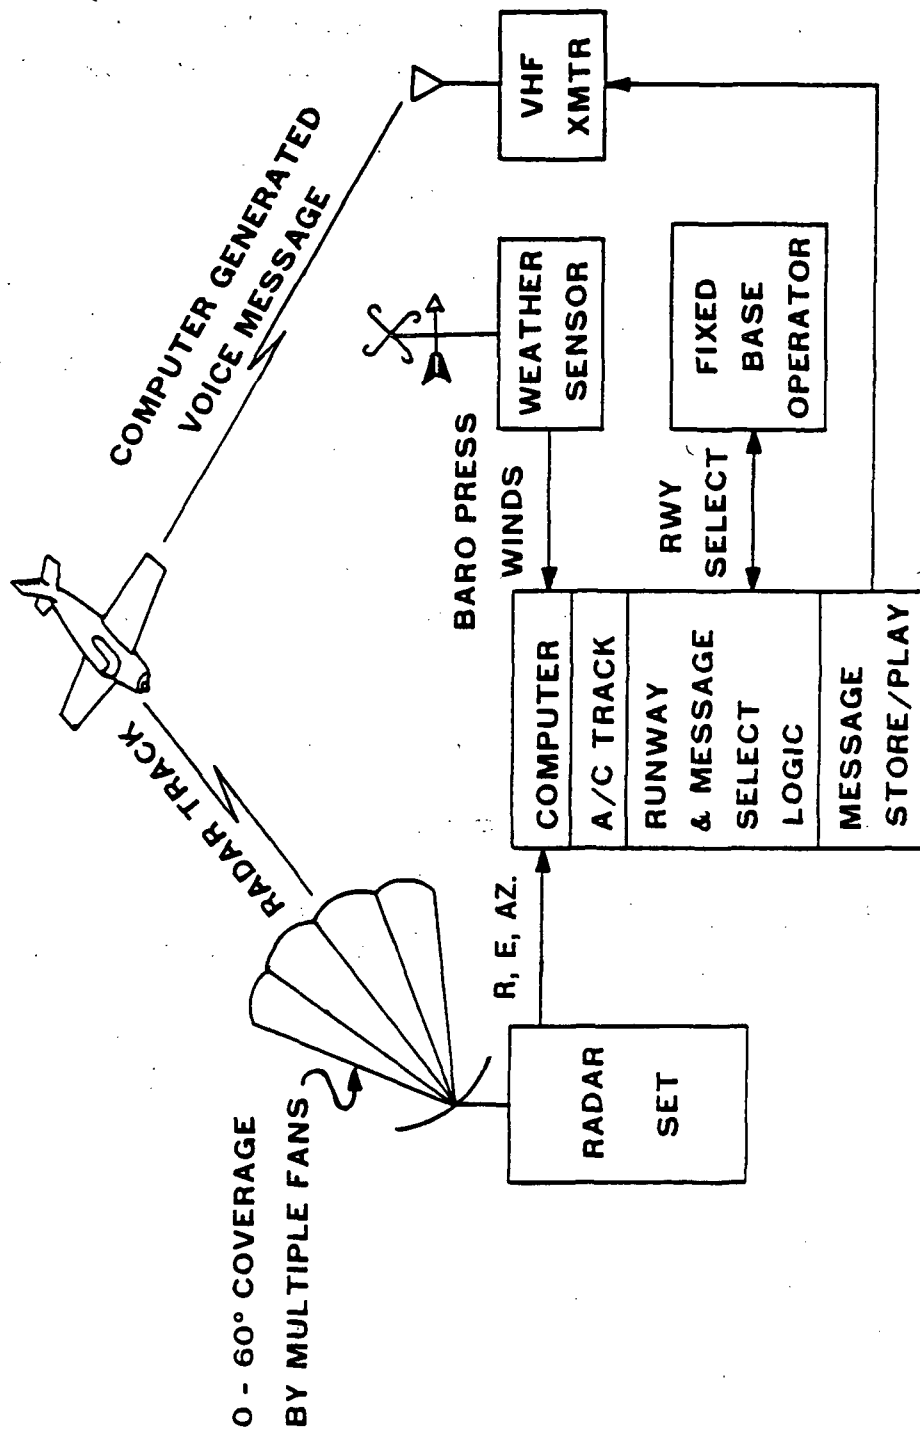


Figure 1-1. - APAS concept.

satisfaction of the following design requirements:

- (1) Low Cost - The system must be affordable to most of the county, municipal, and privately-owned airports in the nation. (A cost limit of \$50,000 in 1975 dollars was imposed for the APAS).
- (2) Airport Advisory System - This system should be capable of:
 - (a) Issuing a report every two minutes at least which would include an airport identifier; time of day; favored runway; wind speed, direction, and gust; altimeter setting; and ambient and dewpoint temperatures;
 - (b) Automatically selecting a preferred runway;
 - (c) Self-testing of weather sensors;
 - (d) Manual control over automatic runway selection and sensor fault declaration via an operator control panel; and
 - (e) Handling at least five additional weather sensors.
- (3) Traffic Advisory System - This system should be capable of:
 - (a) Issuing a report every twenty seconds to identify the number of aircraft on each pattern leg and the range, bearing, and heading of non-pattern aircraft;
 - (b) Radar surveillance of non-cooperative aircraft via a skin tracking radar;
 - (c) Radar coverage to five nautical miles;
 - (d) Height detection;
 - (e) Reporting at least ten aircraft and tracking at least twenty aircraft simultaneously.
- (4) Interface - The APAS should require of pilots only a standard very high frequency (VHF) radio for its use.

As an airport advisory system, APAS is basically an automatic weather station equipped with an automatic voice response unit. Its unique distinguishing characteristic in this respect is its capability of choosing, on the basis of observed wind conditions, a preferred runway for aircraft operations. As a traffic advisory system, APAS becomes what is commonly called an automatic track-while-scan radar system, again equipped with a voice response unit.

The most stringent of the stated requirements concern, as might be expected, is the radar subsystem. If a major goal of APAS is to improve air safety, then clearly a premium is placed on the dependability of traffic advisories, implying, in turn, reliable and accurate tracking of aircraft. It is noted that transponders will not be required of aircraft, since many are not so equipped and interference with normal ATC tracking could occur if APAS were to interrogate these units.

An earlier feasibility study of the APAS concept (Figure 1-1) suggests that the ideal radar for the system would:

- transmit at X-band;
- possess solid-state electronics for reliability;
- perform Doppler processing to alleviate ground clutter problems;
- be capable of detecting a 0.5m^2 target at 3.0 nmi;
- provide 300 m resolution;
- possess a multiple-beam antenna for height finding; and
- be low cost (\$50,000 in production).

The economic feasibility of coherent radars - necessary for Doppler processing - is however, questionable and leads one to favor noncoherent radars with their obvious cost advantages. It is concluded in the feasibility study that sophisticated clutter processing performed in the system computer could compensate for the lack of Doppler information and would render the noncoherent radar technically feasible at all but a small class of problem airports. The present report documents the soundness of this conclusion.

1.2 Outline

This report contains four major sections, each devoted to a different aspect of the experimental APAS:

- System Description - what the system consists of, its operation from the user standpoint, and its approximate cost;
- System Structure - a description of the system in terms of its functional units, their algorithms, and their coordination;
- System Evaluation - how well the experimental system performs and pilot response to it;
- System Development - suggested improvements and alternative approaches.

For the most part, the discussion is kept on an overview level; in particular, there is very little of what could be termed true system documentation - schematics, program listings, etc. Details of this sort are relegated to a companion volume in an effort to make this volume more generally accessible.

In keeping with this general philosophy, the sections which follow are each introduced with a rather extensive overview of the contents of that section. Simply reading these introductions should provide a good knowledge of APAS; the reader can, at his discretion, pursue topics of particular interest by proceeding beyond the introduction.

Sections 2.0, "Description of Experimental APAS," and 4.0, "Test and Evaluation of the Experimental APAS," are perhaps of most immediate interest. Taken together, these sections provide a good discussion of the system from the user point of view: what the system does and how well it does it. It is seen that the experimental APAS consists largely of off-the-shelf equipment carrying a modest price tag. Although it is highly unlikely that a production APAS would use components identical to the experimental version, the fact that equipment in the latter cost approximately \$120,000.00 argues well for economic feasibility of the concept. Of this cost, moreover, a significant fraction (roughly 60%) pertains to computers, their peripherals, or equipment necessary for experimentation. Since many of the peripherals function solely as software development aids and would not appear in a production system, and since the cost of digital hardware continues to decline, it is likely that most of the expense of an APAS would lie in its radar, its weather sensors, and its installation at an airport.

Major components of the experimental system are the following:

- a Marine Pathfinder radar operating at X-band with a 1.0 microsecond pulsewidth and 20 kilowatts transmitted power;

- a single-transmit/multiple-receive antenna scanning 360 degrees every two seconds and step-scanning in elevation;
- an Eclipse S/130 minicomputer equipped with a radar interface processing 4096 radar video samples per second;
- an SBC 80/204 microcomputer equipped with a voice response card;
- an operator control panel providing manual control of the system;
- weather sensors.

A block diagram of the system identifying these major components and their interconnection is shown in Figure 1-2.

The system broadcasts--airport advisories and traffic advisories--occur in an inter-laced fashion, an airport advisory every two minutes and, in between them, traffic advisories every twenty seconds (traffic density permitting). Example broadcasts are as follows:*

AIRPORT ADVISORY. MANASSAS. GEE-EM-TEE ONE THREE FOUR FIVE. FAVORED RUNWAY ONE SIX CHANGING TO THREE FOUR. WIND TWO FIVE ZERO AT ONE FIVE GUSTING TO TWO TWO. ALTIMETER THREE ZERO ZERO FOUR. TEMPERATURE FIVE EIGHT. DEW POINT FOUR THREE. CAUTION - MOWING OPERATIONS IN PROGRESS.

"TRAFFIC ADVISORY. ONE AIRCRAFT ON FINAL. AIRCRAFT ZERO POINT FIVE MILES NORTH HEADING NORTHEAST. AIRCRAFT TWO MILES SOUTH HEADING EAST."

For the most part, these advisories are received favorably by pilots. Users of the system at Manassas airport, asked to respond to a questionnaire in which they could evaluate the system, overwhelmingly preferred APAS traffic advisories over self-announcement (87.5% to 12.5%). In more objective terms, it was found that individual traffic reports broadcast by the system were accurate 95% of the time with no degradation in high traffic density situations. The only serious, often heard complaint about the system concerned its automatic runway selection, an area in which improvement is clearly called for. Section 3.0, "Structure of Experimental APAS," provides a detailed description, albeit at a functional level, of the three major units comprising APAS:

*In APAS phraseology, the period is used as a description inferring end of phrase and/or sentence. A pause between this word and the next succeeding word is inferred.

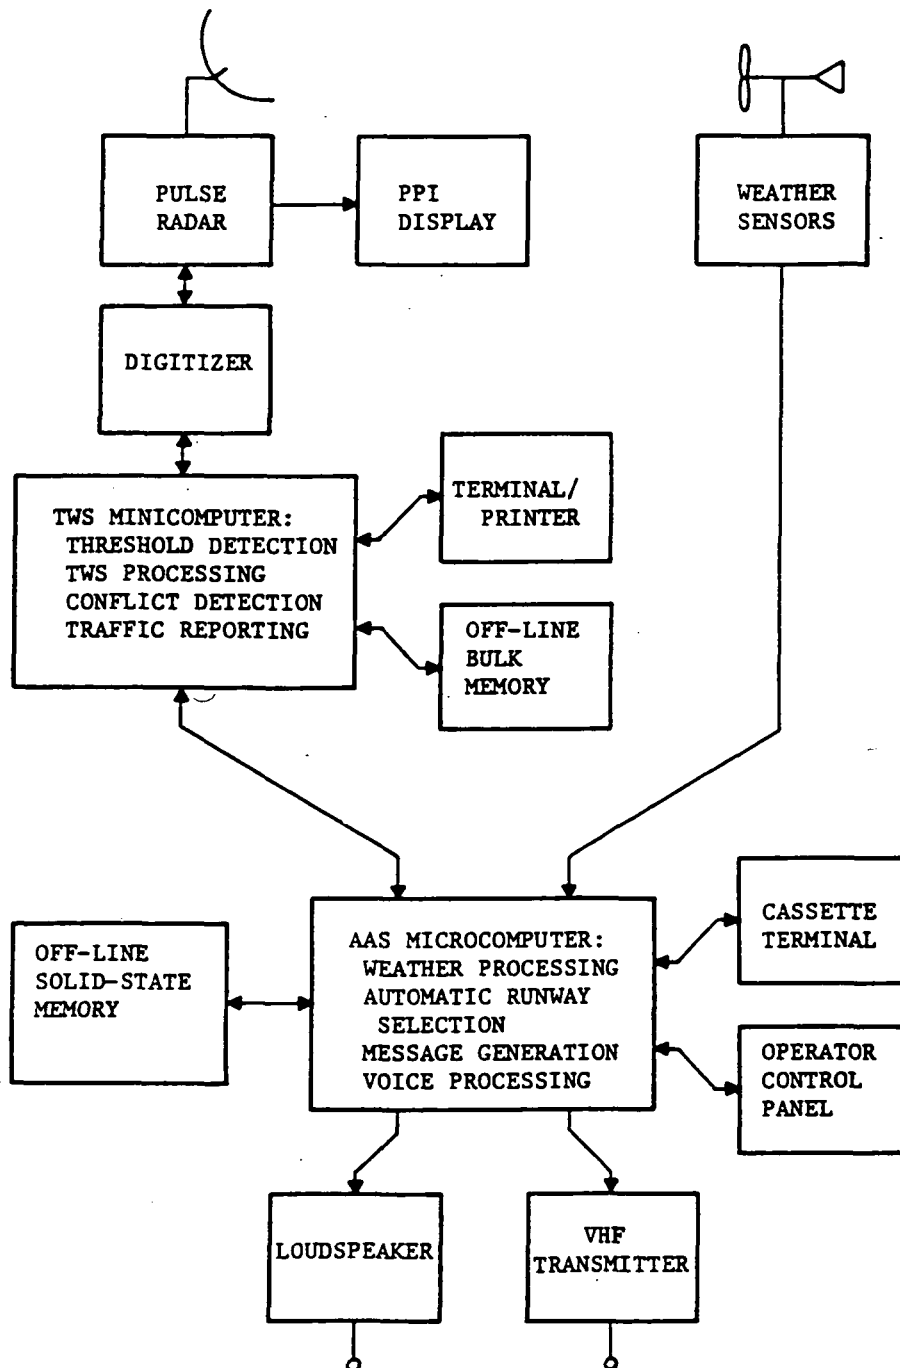


Figure 1-2. - APAS block diagram.

- Tracking Data Unit;
- Weather Data Unit; and
- Voice Response Unit.

The first two of these functional units are the system sensors; the data they provide to the last unit are converted by it into broadcast messages. Each of these units is further subdivided into functional modules as follows:

- Tracking Data Unit
 - Radar Module
 - Target Detection Module
 - Track-While-Scan Module;
- Weather Data Unit
 - Weather Sensing, Processing, and Display Module
 - Runway Selection and Display Module;
- Voice Response Unit
 - Speech Processing Module
 - Message Formatting Module
 - Message Entry Module.

These various modules represent a partitioning of the activities to be performed by the systems; for the most part, they operate semi-autonomously, communicating in a well-defined way with one another. For example, the target detection and track-while-scan modules are separate tasks in the multi-tasking environment of the system minicomputer, communicating with one another through common data blocks and the task scheduler.

Of the three functional units, the tracking data unit is the most complex and warrants the greatest discussion in Section 3.1. Several original algorithms for detection and tracking are developed in this section, algorithms which have proved their worth in field testing. Indeed, the capabilities displayed at times by the system in combatting clutter and jamming are outstanding. The voice response unit is also relatively complex, not because of its speech coding approach (adaptive differential pulse-code modulation), but because of its message formatting capabilities. In an effort to provide maximum flexibility in an experimental system, the message formatting module is designed to accept the rules governing message formation as input data. By comparison with the other two units, the weather data unit is relatively straightforward, its only novelty lying in automatic weather-sensor fault detection.

The discussion of Section 5.0, "Proposed Development," represents a continuation of Section 4.0, "Test and Evaluation of the Experimental APAS." This section treats the structure of the experimental APAS in hindsight, discussing various enhancements which perhaps would have been included in the system had its design followed a year of testing.

The report closes with a summary in Section 6.0.

2.0 DESCRIPTION OF EXPERIMENTAL APAS

An experimental APAS was developed to evaluate and demonstrate the concept defined in Section 1.1 of this report. The system included five subsystems: (1) a radar, antenna array, and clutter suppression; (2) digitizer; (3) computers; (4) weather instruments; and (5) operator control panel. A description of each of the subsystems and the operation of the experimental APAS is presented in the following text. The reader should note that the experimental APAS was designed to emulate as much as practical an operational unit. The primary differences between the experimental system and an operational one are the system flexibility and additional subsystems incorporated for experimentation and the lack of system redundancy required in an operational system.

2.1 Radar, Antenna, and Clutter Suppression Subsystem

The radar used by the experimental APAS was a Raytheon Mariner's Pathfinder surveillance radar, model RM1220-GXR. This radar set includes a transmitter, receiver, plan position indicator (PPI), antenna, and a fixed speed (30 RPM) pedestal. To achieve APAS requirements, major modifications were required in the antenna and pedestal systems and minor modifications in the sensitivity time control (STC) and pulse repetition frequency (PRF) circuits. The modulator transmitter receiver (MTR) and PPI units were essentially unmodified and no modifications were required in the pulse shaping circuit since the radar had selectable 0.5 and 1.0 microsecond pulse widths.

The antenna supplied with the pathfinder radar had an azimuth beamwidth of 1.4° , an elevation beamwidth of 23° , and a gain of 28 db. Because the APAS required height finding capabilities, this antenna could not be used as a receiver antenna and had to be replaced with an array (Figure 2-1). The elevation and azimuth beamwidth of this antenna did make it suitable as a transmit antenna, and it was used in this capacity.

A provision was made for mounting an array of up to five receive antennas which would be operated one at a time in a sequence determined by the digitizer (Section 2.1.1). Antenna selection was accomplished through a single pole, five throw pin diode switch controlled by TTL logic signals from the digitizer. The transmit and receive antenna signal paths (Figure 2-2) were recombined in the E-plane waveguide circulator which was mounted just above the pedestal's rotary joint. To prevent radar signal returns from being received in the transmit antenna and to protect an RF amplifier in the receive antenna line, isolators were incorporated in each path to essentially make each path one-way.

The low-noise RF amplifier in the receive antenna path was required to overcome the losses incurred by the pin diode and splitting the transmit and receive antennas. The addition of this amplifier required the gain of the IF amplifier in the pathfinder's




Figure 2-1. - Antenna array.

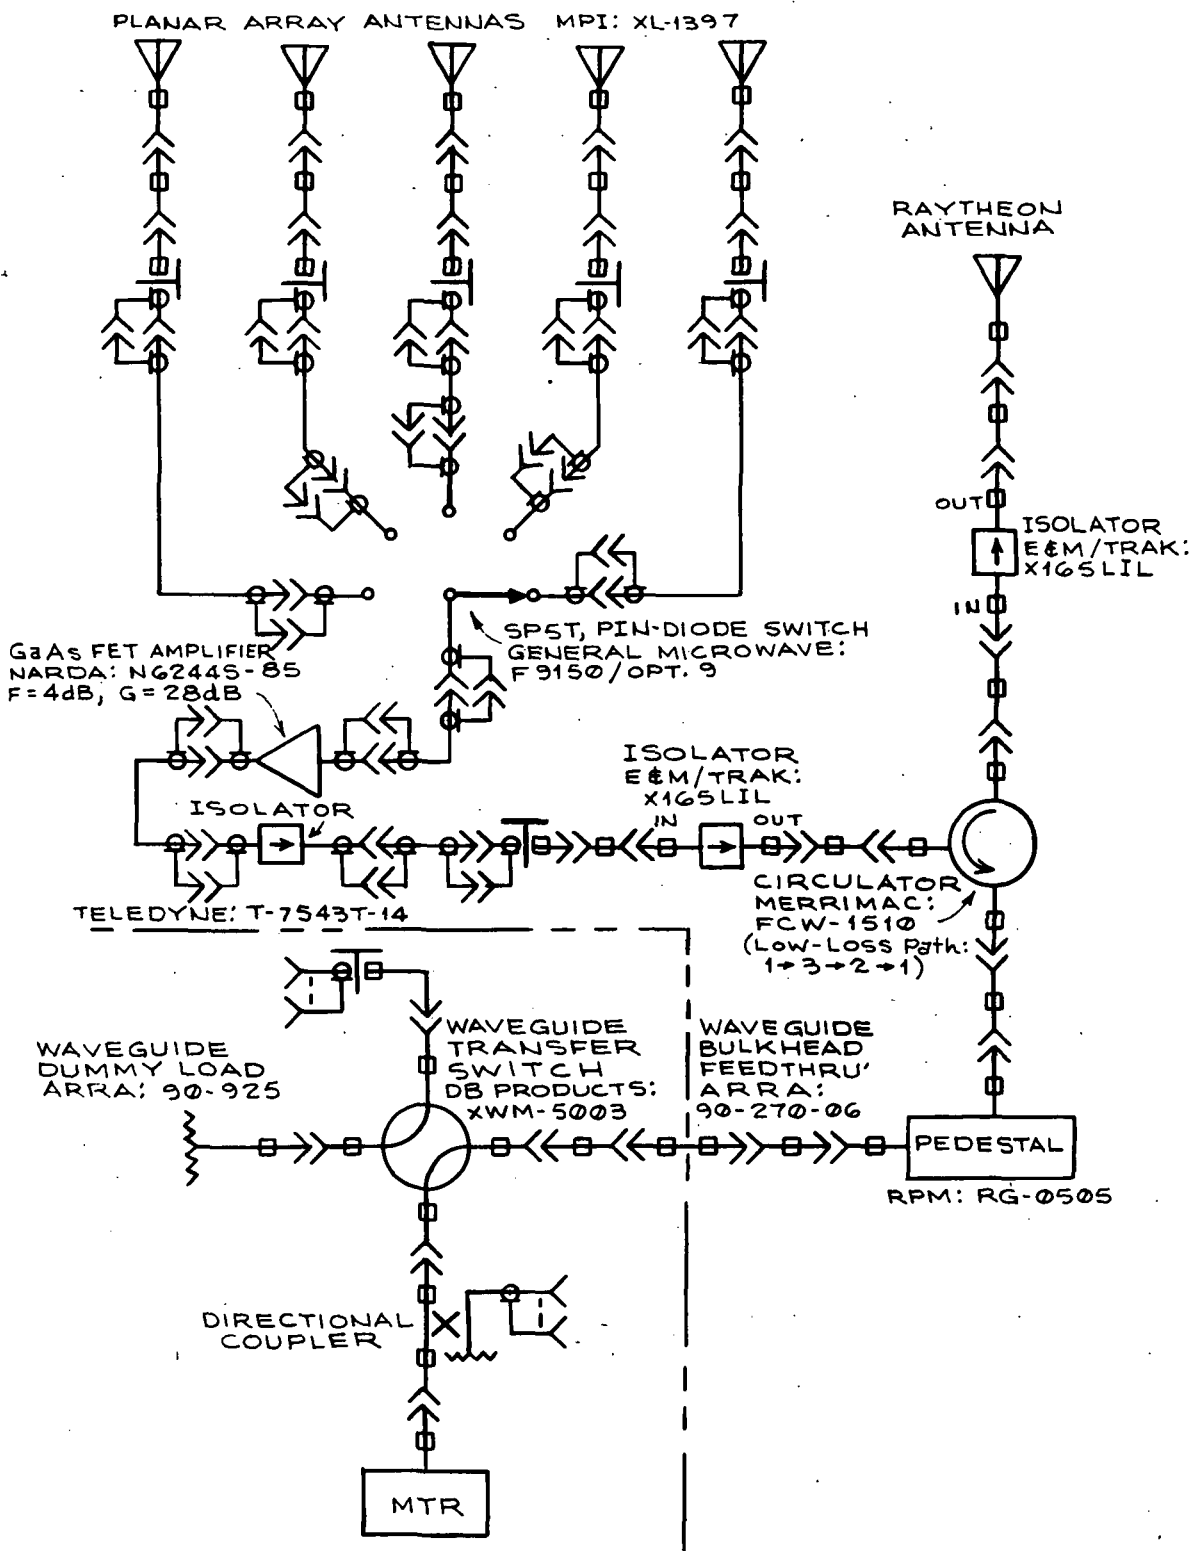


Figure 2-2. - Transmit and receive antenna signal paths.

receiver to be reduced to prevent receiver saturation. This minor modification was accomplished by adding a variable DC bias at the input point for the STC.

The experimental APAS required a variable rate pedestal so that system timing could be optimized. APAS requirements necessitated that the antenna would, in effect, become the system clock, since its role established the rate of data flow, (i.e., "N" number of range information blocks each time the antenna changed 1.4° azimuth sectors). Since the pedestal supplied with the pathfinder was a fixed speed one, a variable speed rotator was procured which allowed speed variations between 1 and 33 RPM. This rotator contained a single channel X-band rotary joint rated at 30 kw peak power, twenty-two DC clip rings rated at five amps each, a one-speed, 60 Hz synchro to provide azimuth data to the digitizer, and a resolver to provide azimuth data compatible with the pathfinder's PPI.

The final modification to the pathfinder radar was to its STC circuit. The STC system was used for clutter suppression and since the APAS used multiple receive antenna set at different elevation angles, an STC circuit was required for each. A new STC unit was designed which provided variable rise time, delay, and bias controls in a circuit (one for each antenna) and whose selection was performed by the digitizer as it selected the receive antenna.

Clutter suppression techniques were required for the APAS because target detection, when using the pathfinder radar, depends on a signal amplitude detection process. References 1 and 2 discuss the various target detection techniques, radar system trade-offs, and the selection of the pathfinder radar for the APAS. For this method to be successful, radar signal returns from ground clutter sources must be suppressed and computer software must be developed to detect and separate the signal returns from aircraft and other sources. (APAS computer software for target detection is discussed in Section 3 of this report).

For the purpose of this report only, clutter signals are defined as all those radar signals received which were not reflected from aircraft (i.e., from trees, buildings, coarse terrain, etc.). Clutter signals would produce two detrimental effects. Since the elevation beamwidths of the receive antenna were somewhat large, radar signals received from aircraft would be mixed with clutter signals at the same range. If the clutter signals were sufficiently large, they would hide the radar signals, and the aircraft would not be detected. The second effect was that clutter signals could produce false target detections unless computer software was designed to detect the difference between clutter and aircraft signals. (Basically, clutter signals occur at the same range over a long period, and aircraft signals occur for a short period.)

The means employed by the APAS to reduce the hiding effect was to employ three systems to reduce the clutter signal: (1) a RF screen; (2) narrow elevation beamwidths

for the lower elevation antenna; and (3) separate STC controls for each receive antenna. A metallic screen wire, with a grade equivalent to that normally used for porch screen, was positioned to reflect both transmitted and received radar signals occurring below a 2° elevation angle from the radar antenna (Figure 2-3). RF signal strength measurements indicated that this wire attenuated the clutter signals by approximately 30 db. The APAS in its final configuration employed three receive antennas to provide elevation coverage. In order to reduce the clutter signal received in each antenna, the elevation beamwidths became a function of the elevation angle of the antenna. The final APAS configuration had antennas at 5° , 10° , and 20° elevations with beamwidths of 3° , 6° , and 13° .

Since radar signals reflected off aircraft at close ranges would be very large compared to threshold detection levels, STC controls were employed on each receive antenna to attenuate the received signals inversely as a function of range. The effect produced was to reduce some clutter signals' returns below levels which could cause the hiding effect while allowing aircraft signal returns to remain at levels at which detection would occur.




Figure 2-3. - RF Screen .

2.1.1 Digitizer

The digitizer is the equipment which provided the interface between the surveillance radar and the tracking computer. Initially, the digitizer was designed to convert the analogue signals from the radar into digital words for inclusion into the tracking computer. As the complexity of an APAS system was realized, it became evident that the digitizer would have to perform additional tasks of computer/radar interface, and it was determined to be the most logical place to institute controls of system performance to achieve experimentation.

Functionally, as an integral component of the APAS, the digitizer performed the following:

1. Converted radar signal returns into digital words.
2. Assembled and transferred to the tracking computer a 64 x 16 bit array of words which included the radar signal returns, the number of the azimuth sector and elevation beam, and error status indicators.
3. Performed error analysis of antenna speed, data transfer, and radar PRF operation.
4. Generated the radar PRF.
5. Converted 1.40 azimuth sectors into 2.8° sectors by either a maximum comparison or the addition of data in each 1.4° sector.
6. Selected the receiver antenna for signal processing.
7. Controlled the selection of the proper STC circuit.

The conversion of radar signal returns into digital words was accomplished by utilizing a high speed A/D converter and 16 x 64 bit memory to sum and store eight integrations of return signals within each 1.4° azimuth sector. Since it would be impossible to slave antenna speed to internal radar system timing, the digitizer was used to externally generate the radar's PRF. The digitizer would produce PRF signals in a burst mode (8 pulses at an equivalent 1300 PRF) each time it detected a 1.4° azimuth sector change. The signal returns from each of the eight pulses would be integrated and stored in one of two memories ("A" memory for the first 1.4° sector and "B" memory for the second 1.4° sector). Upon completion of the 8 pulse integration of "B" memory, the data would be summed or compared and the result stored in a third memory "C".

$$(C_n = A_n + B_n \text{ or } C_n = \max. A_n, B_n.)$$

Since the integrated signal returns of each range bin used less than 15 bits, bits 15 and 16 were used to define the azimuth sector number, elevation beam number, and error status indicators. Error detection circuits were incorporated in the digitizer to verify the proper signal integration, data transfer, and radar PRF operation. An improper operation of either of these circuits would produce an erroneous integrated sum which would cause errors in the clutter mapping and thresholding process.

The digitizer allowed system operation in one of either six operational modes. The two basic modes were single or dual. The dual mode collected and transferred data to the tracking computer in each 1.40 azimuth sector by integration of either 8 or 16 pulses. The single mode collected data in 1.4⁰ azimuth sectors, and by either an addition or comparative process, converted the data into 2.8⁰ azimuth sectors. Both 8 and 16 integrations were permissible in single mode operation.

The digitizer controlled the selection of both the proper STC circuit and receive antenna. This selection was made in an increasing fashion, between antenna number 1 and 5, by stepping to the next highest numbered antenna following a 360⁰ data scan. A selector switch determined the highest numbered antenna for which data was to be taken. Initially, the system was operated with five antenna but later experiments only utilized three.

The digitizer would hold the selected antenna into the highest numbered one (as defined by the number of antenna selector switch) until it received a go from the tracking computer. This feature would allow the tracking computer additional processing time in the event the workload got too heavy. Once the tracking computer determined that it had caught up, it would send signals to the digitizer to proceed. Normally, the proceed signals would be received at 270⁰ azimuth point, and the digitizer would step the antenna select to antenna number 1 at the 360⁰ azimuth point of the last antenna.

2.1.2 Minicomputer and Peripherals

Responsibility for automatic detection and tracking of aircraft targets in APAS is performed by a Data General Eclipse S-130 minicomputer operating under a real-time, disk-based operation system (RDOS). On the input side, the computer is interfaced to the digitizer, accepting integrated radar returns on a per sector basis. It processes these data to create target tracks which form the basis for constructing traffic reports. On the output side, the minicomputer is interfaced to an Intel SBC 80/204 microcomputer which periodically requests from the former such traffic reports. The minicomputer responds to a request by consulting its tracking files for tracks judged to represent aircraft, formatting the pertinent data into a comprehensive traffic report, and communicating the report to the microcomputer for verbalization in a traffic advisory broadcast.

The functions just described, corresponding to APAS being in operational status, are performed by the minicomputer in a multi-tracking environment consisting of three tasks, in order of priority:

- (1) target detection,
- (2) traffic reporting, and
- (3) target tracking.

Responsibility for switching among these three tasks belongs to the task scheduler portion of RDOS. The task of target detection is that with the most demanding data throughput rate and is therefore afforded highest priority as indicated. Although the least demanding of the three tasks, traffic reporting is accorded second priority to insure that the microcomputer not be held up in issuing traffic advisories. Third priority is then given to target tracking by default.

Communication of radar data from the digitizer to the minicomputer occurs in the direct-memory access (DMA) mode of transfer. Thus 64 words of data, corresponding to integrated radar returns from 64 range bins (Section 3.1.1) are output from the digitizer accumulators and stored directly in minicomputer memory. After it has output the last word, the digitizer sends an interrupt to the minicomputer, a signal to the latter that a new set of data is to be processed and that preparations must be made (i.e., a DMA transfer initialized) to receive the next block of 64 words. In contrast, communication between the minicomputer and microcomputer occurs strictly on an interrupt basis, following a standard handshake protocol which uses data available and data received flags.

In addition to software supporting the functions indicated above, the minicomputer is also equipped with various elements of support software, the most important being a set of routines which allow one to create and print out a radar clutter map of an airport terminal area. Clutter mapping, which is discussed in Section 3.1.2, is one of the key techniques employed in APAS to identify stationary reflectors and to alleviate the effects of distributed clutter such as forests. In order to use this technique in APAS, it is necessary to identify, prior to operation, those regions which should be mapped. The indicated software allows one to do so. Completing the package of support software are two programs which allow one simply to test the minicomputer/digitizer and minicomputer/microcomputer interfaces, respectively, in a manner free of extraneous and potentially confusing side effects.

Befitting its status as part of an experimental system, the S-130 is configured as a general purpose minicomputer system, possessing several peripherals which would probably be absent in an operational system. The complete configuration is listed in Table I. Of the indicated components, one of the I/O peripherals (printer, console, disks) is absolutely required for operation of a production system; once started, the system uses these peripherals only for diagnostic and data recording purposes.

2.1.3 Microcomputer and Peripherals

The microcomputer system in APAS, an Intel SBC 80/204, together with its peripherals, is responsible for two main functions, namely weather data acquisition and processing and message broadcasting. Acting alone, this system is capable of delivering airport advisories; interfacing to the minicomputer provides it with the data source

TABLE I. - MINICOMPUTER HARDWARE/SOFTWARE CONFIGURATION

Manufacturer: Data General Corporation

HARDWARE		SOFTWARE	
Model No.	Description	Model No.	Description
8611-K	S/130 Eclipse with 128 KB MOS Memory and BBU	3359F	Eclipse RDOS
		3371F	Eclipse Fortran IV
8613	Floating Point Processor	3375F	Algol Compiler
		3378F	Basic
6045	10 MB DGC Cartridge Disc Subsystem	3476F	Eclipse Fortran IV Compiler with Hardware Floating Point
6030-B	Dual Diskette	3384F	Eclipse Fortran IV Runtime Libraries
6052	CRT	3511F	Fortran Commercial Subroutine Package
4034D	165 CPS Printer	3388F	Eclipse Algol Runtime Libraries
4075	I/O Interface	3473F	Fortran IV Runtime Library with Hardware Floating Point
4077	Asynchronous Line Controller	3562F	Discrete Fourier Transform Library
4078	EIA Interface	3464F	Real-Time Input/Output System
		3604F	Diagnostic Operating System Eclipse
4079	Real-Time Clock	3605F	Diagnostic Operating System Peripherals
4014	I/O Interface	3721-02F	Eclipse Mapped RTOS
4034	Line Printer Control	3374F	Fortran 5 Compiler
1012P	Cabinet	3387F	Fortran 5 Runtime Library

required to deliver traffic advisories as well. In a sense, therefore, the micro-computer can be viewed as the controlling element of APAS, even though its workload is far less than that of the minicomputer.

Off-the-shelf hardware comprising the microcomputer system is listed in Table II. It is noted that the two 64K RAM boards are used for storage of the digitally-encoded words required in constructing broadcast messages. Through a modified addressing scheme built into the machine, a logical address space of 32K is made available for programs and scratch data storage (instead of the 64K normally available with the SBC 80/204). Sacrificing the indicated 32K storage locations, however, allows a logical address space of up to 512K to be used for data storage, of which 128K are actually implemented and hold speech data.

Speech processing is accomplished by a special purpose board designed specifically for the APAS applications and plugged into the microcomputer backplane. Based on the adaptive differential pulse code modulator (ADPCM) coding of speech, this board operates bilaterally as either an encoder or a decoder of speech, the former capability required for the recording of so-called discretionary messages.

Connected to the microcomputer is an operator control panel through which the microcomputer system interfaces with the system operator. On the one hand, the operator control panel displays weather data and the favored runway and makes available locally through a loudspeaker the messages being broadcast by the system. On the other hand, the operator control panel provides the means whereby the system operator can choose active runways, disable weather sensors, disable system broadcasts, or enter verbally discretionary messages.

The weather data sensors used in APAS, listed in Table III, are standard off-the-shelf equipment. They interface to the analog data acquisition board mentioned in the previous figure and are sampled periodically in time. The samples are appropriately processed by the microcomputer to develop the weather statistics suitable for inclusion in airport advisories. As part of this processing, these data are examined for behavior characteristic of a sensor fault. When a fault is detected, microcomputer operation with respect to that sensor enters a fault processing mode in which data from the sensor are suppressed in APAS broadcasts. This mode persists until cleared through the operator control panel.

Operation of the microcomputer is not under the control of an operating system. Rather the application software in the machine is responsible for coordinating its various activities in a real-time environment. For convenience, however, the system is endowed with a rudimentary editor to ease software development and testing, along with a program module to initialize the system and one to accept input data prior to operation. All programs are stored in RAM.

TABLE II. - MICROCOMPUTER SYSTEM CONFIGURATION

Manufacturer: Intel Corporation

Model Number	Description
SBC 80/204	Single Board Computer with 4K RAM
SBC-660	System 80 Chassis with Power Supply
SBC-116	Combination Memory and I/O Expansion Board
SBC-064(2)	64K RAM Boards
SBC-711	Analog Input Board
SBC-905	Universal Prototype Board (Speech Processing Board)
733ASR	Dual Cassette Terminal (Texas Instruments)

TABLE III. - WEATHER SENSOR PACKAGE

Manufacturer: Meteorology Research, Inc.

Model Number	Description
1074-2	Wind Sensor, 540° Azimuth, Light Chopper
751	Barometric Pressure Sensor, 28.0 to 32.0" Hg
892-1	Power Aspirated Ambient and Dew Point Temperature Sensor
1002	Transmuter (4-card capacity), 115 VAC, 300 ma
12905-X	Wind Speed Amplifier, 0 to 75 Knots
14303	Wind Direction Amplifier, 0 to 540° Azimuth
17060	Baro Pressure Amplifier
13495-11	Dew Point Temperature Amplifier, -22° to +122° F.
13495-12	Ambient Temperature Amplifier, -22° to +122° F.
14036	Signal Cable for 1074-2 Wind Sensor to Transmuter
16491	Aspirator Motor Power Cable for 892-1
14714	Baro Pressure Signal Cable to Transmuter
18332	Dew Point Temperature Signal Cable to Transmuter
18976	Ambient Temperature Signal Cable to Transmuter
	Lightning Inhibitors for Line and Amplifiers

2.2 System Operation

To reiterate an earlier statement, APAS is predicated on the hypothesis that providing pilots with timely airport and traffic information will lead to improved safety at high-density, uncontrolled airports. The system disseminates the information in question in the form of broadcast messages of two types: airport advisories, which specify the active runway and weather conditions, and traffic advisories, which describe the existing traffic pattern and generally indicate the position of aircraft flying in the airport terminal area. To receive these advisories a pilot need only tune his VHF receiver to the appropriate frequency.

In routine operation, airport and traffic advisories are broadcast on a regular schedule. Although APAS allows a great deal of flexibility in fixing this schedule, two specific modes of operation have proven satisfactory in practice. The first, representing what is termed earlier the automatic airport advisory system, consists of broadcasting airport advisories alone at a rate of one every twenty seconds. The second scheduling mode, which realizes the complete APAS, interweaves airport and traffic advisories in such a way that the former are broadcast nominally every two minutes and the latter occur at a rate of one every twenty seconds between successive airport advisories. It is possible, however, for the description of a dense traffic situation to require more than twenty seconds, in which case traffic advisories will appear on virtually a continuous basis between airport advisories.

At any time, a fixed-based operator (FBO) can interrupt the normal broadcast schedule to enter into APAS a discretionary message to be appended to airport advisories. When this occurs, the scheduled broadcasts are suspended for approximately twenty seconds.

The sections which follow describe in some detail the format of airport and traffic advisories and the information contained therein. Proceeding in this way allows a convenient and natural description of the operation of APAS as well, in general, with reference to the system user (the GA pilot) but occasionally with reference to the system operator, in practice, the FBO. Section 3 contains any substantive discussion of the means whereby the data appearing in advisories is obtained, the emphasis here being the interpretation to be attached to the data.

2.2.1 Airport Advisories

Airport advisories may be thought of as abridged versions of those messages broadcast by the automatic terminal information service (ATIS). In the case of APAS, message content is modified to conform with the data available from an unattended system (e.g., reported ceilings are not announced automatically) and to suit the needs of an uncontrolled airport.

Information included in an airport advisory can be classified as follows:

- 2.2.1.1 Airfield name
- 2.2.1.2 Time-of-day
- 2.2.1.3 Wind conditions:
 - Direction
 - Speed
 - Gusts
- 2.2.1.4 Runway status:
 - Current active runway
 - Next active runway
- 2.2.1.5 Altimeter
- 2.2.1.6 Temperature and dew point
- 2.2.1.7 FBO message (discretionary message)

An illustrative, if somewhat fanciful, advisory which contains references to data in each of these categories is the following:

AIRPORT ADVISORY. MANASSAS. GEE-EM-TEE ONE THREE FOUR FIVE. FAVORED RUNWAY ONE SIX CHANGING TO THREE FOUR. WIND TWO FIVE ZERO AT ONE FIVE GUSTING TO TWO TWO. ALTIMETER THREE ZERO ZERO FOUR. TEMPERATURE FIVE EIGHT. DEW POINT FOUR THREE. CAUTION - MOVING OPERATIONS IN PROGRESS.

Depending on various data conditions, the exact form of an airport advisory may vary considerably. The rules governing message construction are detailed in the following subsections.

2.2.1.1 Airfield name. Following the header "AIRPORT ADVISORY." is announced the name of the airfield at which APAS is located. In the illustrative airport advisory above, "MANASSAS." identified Manassas Municipal Airport in Manassas, Virginia. If no airfield identifier is indicated to the system at start-up, the airfield name announcement is omitted from airport advisories.

2.2.1.2 Time-of-day. Greenwich Mean Time (GMT) is announced in every airport advisory if the system clock is set by the FBO at start-up and is omitted from all advisories otherwise. Time is always specified by the four digits of the twenty-four hour clock, so that "GEE-EM-TEE ONE THREE FOUR FIVE." indicates a GMT of 1345 hours.

2.2.1.3 Wind conditions. In common with its other meteorological sensors, APAS acquires samples from the wind sensors (direction and speed) at four-second intervals. Sample storage is arranged so that, at any point in time, the system can compute the following three wind statistics:

- (1) average wind direction during the preceding one minute,
- (2) average wind speed during the same time interval, and
- (3) peak wind speed during the preceding five minutes.

Unless specified otherwise in what follows, these summary statistics are referred to simply as wind direction, wind speed, and wind gusts, there being little need to differentiate the first two explicitly from the underlying raw data samples.

In addition to forming an important part of airport advisories, the wind statistics (speed and direction) serve two other functions: input to the runway selection process and input to the track-while-scan process. In the first instance, wind data plays the crucial role in identifying the preferred runway to serve as the active runway. In the second case, wind data have proven a valuable aid in tracking aircraft, particularly aircraft in the traffic pattern.

As far as their inclusion in airport advisories is concerned, the wind statistics are used to prepare a statement of wind conditions whose format may vary considerably, depending primarily on wind speed. If, for the moment, it is assumed that both wind sensors are operational, then several cases can be distinguished according to the following outline:

Calm winds. Wind speed less than four knots.

Moderate or steady winds. Wind speed greater than four knots and wind gusts either less than fifteen knots or less than five knots above wind speed.

Strong, gusting winds. Wind speed greater than four knots and wind gusts both greater than fifteen knots and greater than five knots above wind speed.

In the first case of calm winds, the wind announcement takes the abbreviated form: "WINDS LIGHT AND VARIABLE.". In the case of moderate or steady winds, the wind announcement explicitly mentions both direction and speed: "WIND TWO FIVE ZERO AT ONE FIVE."

indicating a wind speed of fifteen knots and a wind direction of approximately 250°. (The wind direction is always rounded off to the nearest multiple of 10°.) For the final case of strong, gusting winds, a reference to wind gusts is appended to the statement of the last example: "WIND TWO FIVE ZERO AT ONE FIVE GUSTING TO TWO TWO." To reiterate, all of these illustrative wind announcements apply to the case in which both wind sensors are operational.

It is possible, however, for any weather sensor, and in particular the wind sensors, to enter a non-operational state either through the action of the FBO in declaring a sensor out of order on the operator control panel or through a decision by the system that it is receiving ill-behaved data from a sensor indicative of a malfunction. If the wind direction sensor is declared non-operational, then the wind announcement becomes: "WIND NOT AVAILABLE.". In this case, automatic runway selection is terminated (see below). If only the wind speed sensor is non-operational, however, the wind announcement continues to refer to direction: "WIND TWO FIVE ZERO.". Automatic runway selection is retained in this case by invoking the assumption that the (unknown) wind speed is a constant five knots (see below).

2.2.1.4 Runway status. In choosing an active runway, APAS bases its selection on four considerations:

- (1) those runways deemed suitable by the FBO to serve as the active runway (as indicated by switch settings on the operator control panel),
- (2) preferential ordering of runways specified by the FBO to the system at start-up,
- (3) the current wind conditions (direction and speed), and
- (4) the current active runway.

Generally speaking, the selection algorithm (discussed in detail in Section 3.2.3) determines, from the available runway candidates as specified by the FBO, that runway most ideally suited to the prevailing wind. The selection reflects both the relative desirability of the various runways (preference could be given, for example, to paved strips) and the current runway situation (to prevent overly frequency changes in the active runway).

The system invokes the runway selection algorithm periodically to determine if the current runway is still suitable for the prevailing wind conditions. If so, then the runway announcement makes reference to this runway alone; for example, "ACTIVE RUNWAY ONE SIX.". If runway 16 happens to call for a right-hand approach, then such is indicated in the announcement: "ACTIVE RUNWAY ONE SIX RIGHT-HAND PATTERN.". If, on the other hand, a change in the current active runway is indicated, then the runway announcement indicates the impending change: "ACTIVE RUNWAY ONE SIX CHANGING TO THREE FOUR." In this case, the right-hand pattern modifier, if appropriate, is applied only to the next active runway, in this case, runway 34.

The selection of an active runway and the broadcast of airport and traffic advisories is coordinated so that a minimum of one minute elapses between a runway announcement indicating a change, say to runway 34, and the announcement indicating that runway 34 is the active one, thereby affording pilots the opportunity to adjust to a changing runway. Aircraft on final, for example, can complete their approach; those on downwind can execute a go-around. With reference to the two broadcasting modes described earlier, the following would occur during a runway change: If airport advisories alone were being broadcast every twenty seconds, then the runway change would be announced in three successive advisories. If airport advisories and traffic advisories both were being broadcast, the former every two minutes, then the runway change announcement would occur in one airport advisory and in all succeeding traffic advisories (see below) until the next airport advisory broadcast.

The precise wording used for the runway status announcement varies according to whether the system is in automatic or manual runway-selection mode. These two modes are distinguished by the number of runway candidates indicated on the operator control panel. If more than one runway switch is in the auto position, then the system is, by definition, in automatic runway-selection mode, i.e., it has several runway candidates from which to select the active runway. If only a single switch is in the auto position, however, the system is, by default, in manual mode, i.e., it has but one candidate for the active runway, that selected manually by the FBO. In both cases, the runway selection procedure per se can and does remain the same. The results of runway selection are announced differently, however, with the terms "FAVORED" and "ACTIVE" being used for automatic mode and manual mode, respectively. With respect to the first example above, one has "FAVORED RUNWAY ONE SIX." if the system is in automatic mode.

In the case of wind sensor faults, one of several possibilities exist for runway status announcement. If the system is in manual mode, there is no change in the announcement, the FBO is overriding the system anyway, and wind conditions are irrelevant to runway selection. If the system is in automatic mode, then the wind announcement, "WIND NOT AVAILABLE.", corresponding to a non-operational wind direction sensor, causes the runway announcement to be deleted entirely. If the wind direction sensor remains operational, however, automatic runway selection is allowed to continue under the assumption of mild wind conditions, i.e., a constant wind speed of five knots.

2.2.1.5 Altimeter. From the four-second barometric pressure samples, APAS selects the most recent to convert into an altimeter reading for broadcast. This particular sample is subjected to standard NASA procedures (see Section 3.2.2) to correct for airport altitude, a parameter which is entered into the system at start-up. The resulting altimeter announcement takes the form: "ALTIMETER THREE ZERO ZERO FOUR.", indicating 30.04 inches of Hg. In the case of a non-operational sensor, the altimeter announcement follows the form of the wind announcement in similar circumstances: "ALTIMETER NOT

AVAILABLE."

2.2.1.6 Temperature and dewpoint. With only minor differences, temperature and dew point data are acquired, processed, and announced in the same way as altimeter data: "TEMPERATURE FIVE EIGHT. DEW POINT FOUR THREE.", indicating a dry-bulb temperature of 58⁰ Fahrenheit and a wet-bulb temperature of 43⁰ Fahrenheit. In both instances, the most recent data sample is announced, but with no correction as is done for barometric pressure data. In either case, a non-operational sensor causes the standard announcement; e.g., "TEMPERATURE NOT AVAILABLE."

2.2.1.7 Discretionary message. All of the announcements described above are constructed from a vocabulary of words which are entered into the system at start-up and are thenceforth permanently resident in the system. By contrast, the discretionary message announcement is entered orally by the FBO into the system after start-up and may be changed or suppressed at his discretion at any time without restarting the system. In essence, APAS simply records the FBO's voice and plays it back at the appropriate time. Subject only to the constraint that it be less than five seconds long, the discretionary message announcement may take any form whatsoever.

To enter a discretionary message into the system, the FBO depresses the record switch on the operator control panel. Doing so causes the red acknowledgement light to come on, indicating to the FBO that the system is aware of his intentions. After a delay during which pending airport and traffic advisories (at most one of each) are broadcast, the system indicates its readiness to accept the discretionary message by extinguishing the red light. After a short delay of roughly one second, the green record light comes on indicating that the system is recording the signal being received on the control panel microphone and that the FBO should begin speaking. Recording occurs as long as the green light is on (roughly five seconds) at the end of which time the system replays the just-entered message for the sole benefit of the FBO (it is not broadcast).

At this time, the system interrupts its normal schedule to broadcast immediately an airport advisory, usually with the new discretionary message appended to it. The FBO may, however, suppress broadcast of the message by simply setting the message enable switch to the off position. At the cost, therefore, of a minor one-time change in the usual broadcast schedule, a discretionary message may be kept resident in the system indefinitely and yet broadcast only at selected times during the day.

2.2.2 Traffic Advisories

Messages of this type are summary descriptions of the air traffic conditions in the airspace surrounding an uncontrolled airport. Traffic advisories are intended primarily to make pilots aware of other aircraft, particularly nearby aircraft or those with which a pilot can reasonably anticipate interacting. With this information, a pilot can better

plan his approach to, or departure from, the airport and, it is hoped, see and avoid aircraft in a potential conflict situation.

The information broadcast in traffic advisories falls into three categories:

- (1) pattern aircraft,
- (2) non-pattern aircraft, and
- (3) runway status changes.

A typical traffic advisory might read as follows: "TRAFFIC ADVISORY. ONE AIRCRAFT ON FINAL. AIRCRAFT ZERO POINT FIVE MILES NORTH HEADING NORTHEAST. AIRCRAFT TWO MILES SOUTH HEADING EAST.". As with airport advisories, the precise form of a traffic advisory is subject to wide variation most notably in the description afforded to each of the aircraft mentioned therein. The subsections which follow treat the rules governing the formation of a traffic advisory.

2.2.2.1 Pattern aircraft. As implied above, APAS divides all aircraft under track into two major categories: (1) pattern aircraft; i.e., those flying an approach pattern, and (2) non-pattern aircraft.* To be classified as a pattern aircraft, an aircraft must satisfy two conditions: it must be near the active runway; i.e., in the pattern area, and its position and heading must be such as to place the aircraft on a pattern leg. Failure to satisfy either of these conditions disqualifies the aircraft from the pattern category. Five pattern legs are recognized by APAS: final, base, downwind, crosswind, and upwind (illustrated in Figure 2-4).

To define the pattern classification procedure succinctly, let (x, y, z) represent the position of an aircraft and $(\dot{x}, \dot{y}, \dot{z})$ its velocity. These sets of coordinates are defined relative to a fixed, left-handed, rectangular coordinate system centered at the radar set with its x-axis pointed north. Since this coordinate system is awkward for pattern classification, a transformation to runway-oriented coordinates is made. If, according to Figure 2-5, (x_R, y_R) denotes the center of the active runway, the position and velocity of the aircraft in a north-oriented coordinate system whose origin is at $(s_R, y_R, 0)$ are given by:

$$x' = x - x_R,$$

$$y' = y - y_R,$$

$$z' = z;$$

*Although not directly pertinent to the discussion at hand, the question of whether a given track does indeed represent an aircraft (and not, say, ground clutter) must of course be addressed before the pattern/non-pattern distinction can be made. This question is considered in Section 3.1.3.




Figure 2-4. - APAS pattern legs.




Figure 2-5. - Runway coordinate system.

and

$$\dot{x}' = \dot{x},$$

$$\dot{y}' = \dot{y},$$

$$\dot{z}' = \dot{z}.$$

A second transformation (rotation) to the runway heading:

$$x'' = \cos(\psi_R) x' + \sin(\psi_R) y',$$

$$y'' = \sin(\psi_R) x' + \cos(\psi_R) y',$$

$$z'' = z',$$

gives the aircraft position in a runway-oriented coordinate system, where ψ_R is the runway heading. The aircraft velocity is likewise given by:

$$\dot{x}'' = \cos(\psi_R) \dot{x}' + \sin(\psi_R) \dot{y}',$$

$$\dot{y}'' = \sin(\psi_R) \dot{x}' + \cos(\psi_R) \dot{y}',$$

$$\dot{z}'' = \dot{z}'.$$

In terms of the newly defined coordinates, the first condition which must be satisfied by a pattern aircraft is that it be near the active runway. Thus one must have*

$$(x'^2 + y'^2)^{1/2} < D_p$$

and

$$z' < Z_p,$$

*The various constraints used here and below to categorize aircraft are written in terms of strict inequalities. In practice, of course, floating-point arithmetic used in minicomputer does not distinguish between $<$ and \leq .

where D_p and Z_p are parameters which specify, with respect to a given airport, a cylinder within which almost all aircraft fly when they are actually in the pattern. The precise values chosen for these parameters depend on empirical observation of the existing traffic pattern and will vary from airport to airport. Aircraft which do not satisfy the first condition are said to be outside the pattern. Those which satisfy this condition, but not the second, are said to be above the pattern, a distinction which is important in discussing non-pattern aircraft below.

Given that an aircraft satisfies the qualifying conditions above, an attempt is next made to place the aircraft on a pattern leg. To be located on a given pattern leg, the aircraft position (x'', y'', z'') and velocity $(\dot{x}'', \dot{y}'', \dot{z}'')$, with respect to the runway, must satisfy a set of conditions unique to that pattern leg. These conditions are described below, wherein (d'', θ'') is the polar coordinate representation of (x'', y'') and (q'', ψ'') that of (\dot{x}'', \dot{y}'') . In particular, θ'' is aircraft azimuth and ψ'' its heading with respect to the runway centerline.

A preliminary step in the classification process is the identification of so-called pattern departures; i.e., departing aircraft which have not yet cleared the pattern area. This class of non-pattern aircraft, whose description in a traffic advisory is treated in the next section, are defined by the following conditions (see Figure 2-6):

$$\psi'' < 40^\circ \text{ or } \psi'' > 328^\circ$$

and

$$x'' > \frac{L_R}{2}$$

and

$$|y''| < 500 \text{ m or } \theta'' < 15^\circ \text{ or } \theta'' > 345^\circ,$$

where L_R is the length of the active runway.

Having dispensed with the pattern departures, the remaining steps of the pattern classification process involve a check of the following possibilities:

- Upwind Leg:

$$\psi'' < 40^\circ \text{ or } \psi'' > 328^\circ$$

and




Figure 2-6. - Departing aircraft coordinate system.

$$x'' > - \frac{L_R}{2}$$

and

$$y'' > 250 \text{ m.}$$

- Crosswind Leg:

$$238^0 < \psi'' < 328^0$$

and

$$x > \frac{L_R}{4}$$

- Downwind Leg:

$$148^0 < \psi'' < 238^0$$

and

$$y'' < - \frac{W_R}{2},$$

where W_R is the width of the active runway.

- Base Leg:

$$58^0 < \psi'' < 148^0$$

and

$$x'' < - \frac{L_R}{2}$$

and

$$y'' < 950 \text{ m.}$$

o Final Leg:

$$\psi'' < 58^{\circ} \text{ or } \psi'' > 328^{\circ}$$

and

$$x'' < - \frac{L_R}{2}$$

and

$$|y''| < 950 \text{ m}$$

and

$$\frac{z''}{d''} < \tan(23^{\circ}).$$

These various sets of conditions are illustrated in Figure 2-7.

The results of pattern classification, as expressed in traffic advisories, is a series of counts, each count representing the number of aircraft on the respective pattern leg. If a given count is non-zero, then the count is verbalized and identified by the corresponding pattern leg. Otherwise, the pattern leg is not mentioned. Thus, for example, one could have: "ONE AIRCRAFT ON FINAL. TWO AIRCRAFT ON DOWNWIND. ONE AIRCRAFT ON UPWIND.". On the other hand, if all counts are zero, then no mention is made of pattern legs whatsoever.

Because of delays inherent in tracking and classifying aircraft, it is possible that an aircraft reported on, say, downwind, has in fact begun or even completed its turn to base. In view of this circumstance, the pattern leg counts should be interpreted not literally but rather as indicative of the total number of pattern aircraft and their distribution about the pattern as it existed at some earlier moment, perhaps ten to twenty seconds in the past.

An additional pattern class, only partially implemented in APAS but to be verbalized as those just described, concerns aircraft on or over the active runway. This class is defined by the conditions:

$$\psi'' < 40^{\circ} \text{ or } \psi'' > 328^{\circ}$$

and





Figure 2-7. - Pattern leg coordinate system.

$$|x| < \frac{L_R}{2}$$

and

$$|y| < \frac{W_R}{2}$$

Since it is virtually impossible to detect whether the aircraft is indeed on the ground or not, the less restrictive description "OVER RUNWAY" is recommended for this pattern class. Thus one could have as part of the pattern aircraft announcements: "ONE AIRCRAFT OVER RUNWAY."

Needless to say, the entire discussion of pattern aircraft to this point is with reference to a left-hand pattern. The same approach, however, is immediately applicable to a right-hand pattern if one simply replaces y and \dot{y} in the preceding conditions with their negatives. The geometric interpretation of this change of sign is to reflect the diagrams of Figure 2-7 about the extended centerline of the active runway.

In closing this section, it is noted that APAS does not require aircraft to fly each of the five pattern legs recognized by the system. Although, in tracking aircraft, the system may resort from time-to-time on the overall structure imposed on flight paths by the existence of a pattern, its successful operation does not depend on pilots following rigorously defined trajectories. On the basis of available data, APAS simply attempts to infer whether an aircraft is in the pattern and, if so, what pattern leg it is on. The correctness of any decision in this regard is to be judged more on whether a given aircraft can thereby be located successfully by other pilots and less on whether the broadcast information truly reflects the intentions of the pilot of the aircraft in question.

2.2.2.2 Non-pattern aircraft. All aircraft which are not identified as being in the pattern are, by default, considered non-pattern aircraft. In order to merit mention in a traffic advisory, however, it is necessary that such an aircraft not be an overflight. Thus, with reference to the symbology of the previous section, it is required that the altitude of the aircraft satisfy:

$$z' < Z_T,$$

where Z_T defines the maximum altitude of an overflight (nominally 3,000 ft.). Violation of this constraint causes the aircraft in question to be ignored. Each aircraft passing the overflight test, on the other hand, is accorded a full sentence

description in each traffic advisory. In its core form, this description includes aircraft heading as well as range and bearing from the active runway.

If, then, the position and velocity of an aircraft are (x', y', z') and $(\dot{x}, \dot{y}, \dot{z})$, respectively, its range and bearing with respect to the center of the active runway are d' and θ' , respectively, where (d', θ') is the polar representation of (x', y') . Likewise, ψ' is its heading, where (q', ψ') is the polar coordinate representation of (x', y') . For verbalization purposes, d' is rounded to the nearest one-half nautical mile (but never to zero) and both θ' and ψ' are rounded to one of the standard eight points of the compass. With respect to the earlier example, one might have, therefore, "AIRCRAFT ZERO POINT FIVE MILES NORTH HEADING NORTHEAST.", indicating an aircraft for which:

$$d' < 0.75 \text{ nmi},$$

$$337.5^\circ < \theta' \text{ or } \theta' < 22.5^\circ$$

and

$$22.5^\circ < \psi' < 67.5^\circ.$$

If it should happen that an aircraft is above the pattern, i.e.,

$$d' < D_p$$

but

$$z' > Z_p,$$

then an appropriate modifier is added to the description. In the sentence above, for example, one might have: "AIRCRAFT ZERO POINT FIVE MILES NORTH HEADING NORTHEAST ABOVE PATTERN ALTITUDE." if $D_p > 0.75$ nmi (typically the case) and if $z' > Z_p$. A second special modifier is used for pattern departures defined in the preceding section, namely the adjective "DEPARTING": "DEPARTING AIRCRAFT ONE POINT ZERO MILES NORTHEAST HEADING NORTHEAST.". In no other case is the word "DEPARTING" (or its opposite, "ARRIVING") used.

In the case of heavy traffic, it is possible that a traffic advisory may grow overly long if there are many non-pattern aircraft. For this reason, ten non-pattern aircraft at most are described in any given traffic advisory, specifically those ten closest to the runway. The order in which the aircraft are reported follows the points

of the compass, beginning with north; aircraft within the same compass sector are reported in order of increasing range (r').

At the opposite extreme is the case in which there are no aircraft to be mentioned in a traffic advisory. In this special case, an advisory will take the abbreviated form: "TRAFFIC ADVISORY. ZERO AIRCRAFT."

2.2.2.3 Runway Status. In airport advisories, it will be recalled, a runway announcement is included except under very unusual circumstances. By contrast, such an announcement is made in traffic advisories only when a change in the active runway is in progress. Thus the announcement of a runway change in an airport advisory triggers the same announcement in all traffic advisories until the change is finalized. Given that several traffic advisories will be broadcast between successive airport advisories, this approach allows pilots to be alerted several times of a pending runway change but does not bother them with repetitive descriptions of a static runway situation.

Furthermore, a pending runway change is, with respect to traffic advisories, treated as if there were no active runway, i.e., pattern aircraft, described in a preceding section, no longer exist: all aircraft are considered not to be flying a pattern, even though some may in fact be completing their approach to the current, about to be changed, active runway.

3.0 STRUCTURE OF EXPERIMENTAL APAS

To create and broadcast the airport and traffic advisory messages, the experimental APAS was structured around three functional units: (1) tracking data unit; (2) weather data unit; and (3) the voice response unit. Each of these units are primarily computer software and are required to coordinate their activities with the other units. This coordination, and a description of the algorithms of each, are presented in the preceding sections.

Section 3.1 describes the tracking data unit. The information presented in this section describes the methods employed to process the radar data, to perform target detection, tracking, classification, and report validated tracks.

Section 3.2 describes the weather data unit. This unit is responsible for processing and validating the weather sensory data, reporting and displaying the data, and selecting the favored runway from the candidate runways.

Section 3.3 describes the voice response unit which is responsible for creating the voice messages to announce the tracking and weather sensory data. In addition to computer software, this unit contains the special purpose board (Section 2.1.3) which converts the digital speech into analog signals. These signals are directed to a VHF transmitter which broadcast the two advisory messages.

3.1 Tracking Data Unit

The tracking data unit (TDU) is composed of the three modules as shown in Figure 3-1: the radar module, the target detection module, and the track-while-scan module. Together these modules create what is customarily termed an automatic-detect, automatic-track-while-scan radar system; i.e., a surveillance radar equipped with subsidiary computation equipment which detects and tracks multiple targets automatically.

3.1.1 Radar Module

It is the function of the radar module to perform continuous radar surveillance of the region about an airport and to report the results of this ongoing search to the target detection module. The data supplied to the latter module consist of integrated radar returns which are to be examined for the presence of targets of interest.

The radar module, shown in Figure 3-2, includes those elements commonly associated with a radar system, namely transmitter, receiver, antenna, and PPI display. In APAS, the display is relegated from its usual role as a primary radar output device to one of acting simply as a monitor on overall system operation. In its place as the primary radar output device is a digitizer, discussed in Section 2.1.1, which serves as the critical element in the interface between the radar set and the minicomputer, converting the analog signal from the radar receiver into a digital signal suitable for input to the minicomputer. Also part of this interface is the video data handler in the minicomputer which accepts the incoming data stream, examines it for errors, and eventually writes the data into a buffer memory.

3.1.1.1 Data acquisition. As the foregoing introduction indicates, the primary task of the radar module is one of data acquisition. It will be seen that a fixed-window processing scheme is used in the TDU, wherein the 360° azimuthal coverage of the system is divided into distinct resolution sectors, fixed in space, each of which are processed more or less independently from the presence of targets. The particular approach to data acquisition implemented in the radar module (discussed in the following paragraphs) is predicated on this processing philosophy.

Although the radar module allows the system operator some latitude in the control of system parameters (such as scan speed, pulse repetition frequency, etc.), for the most part the discussion emphasizes those parameter values which have proven satisfactory in practice.

The multiple reflector antenna assembly described in Section 2.1 is operated in a step-scan mode in which the output of only one reflector is routed, at any point in time, to the radar receiver. In this scheme, the entire antenna assembly rotates as a unit, the outputs of all reflectors being fed to a rotary switch. As indicated in Figure 3-1,

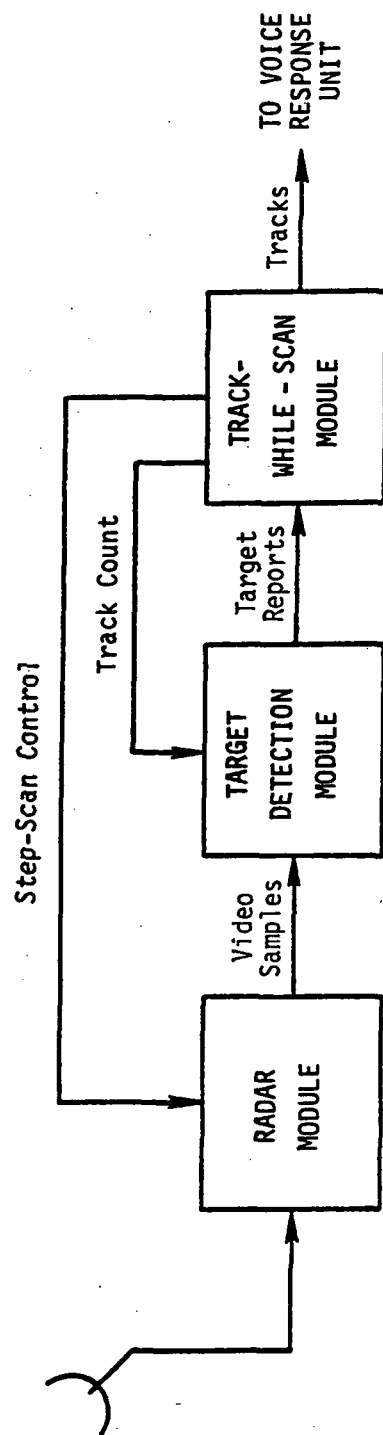


Figure 3-1. - Tracking data unit.

each reflector acquires a numerical designation from its position on the switch; e.g., reflector number 1 is the one connected to switch position number 1. The switch in turn allows the output of a single reflector to appear at the receiver front-end. Orderly stepping is accomplished by slaving the antenna switch to the rotation of the antenna assembly: every time the assembly rotates through a preset reference direction, stepping occurs, proceeding cyclically through the various switch positions in ascending order.

Figure 3-2 indicates that the radar module will accommodate up to $N^0 = 5$ individual reflectors, each with its own sensitivity time control. Fewer than this number may be used, however, provided only that they are assigned to the lowest order switch positions. If N_ϕ reflectors are used, then N_ϕ scans (i.e., complete rotations) of the antenna assembly constitute one complete survey by the system, during which each reflector will have been active for precisely one scan. With a scan time of T_σ seconds, a complete survey requires $N_\phi T_\sigma$ seconds. Stated another way, aircraft should be detected and their tracks updated at least once every $N_\phi T_\sigma$ seconds. In the current APAS configuration, $N_\phi = 3$ and $T_\sigma = 2$ seconds, implying a nominal track update rate of once every 6 seconds.

A subsidiary capability of the step-scan process as implemented in the radar module is beam-holding, a technique for acquiring additional computing time for the target detection and track-while-scan modules in the case of excessive processing backlog. After completing a survey, the radar module simply inhibits antenna stepping until ordered by the track-while-scan module to release it. In this way, additional time in increments of T_σ seconds is made available for detection and tracking at the expense of a longer effective survey time. In normal operation, however, the tracking data unit keeps pace with the incoming data and beam-holding is not required.

As the antenna assembly rotates through a given scan, pulses are regularly transmitted and their returns duly received. In consonance with the fixed-window processing philosophy, the radar module processes precisely n_p pulses per elemental azimuth sector, by which is meant a sector formed by dividing the circle into $N_\theta^0 = 256$ equal parts, each of width 1.4° . Stated in operational terms, the digitizer expects to process at least n_p pulses as the antenna rotates through an angle

$$(j-1) \frac{2\pi}{N_\theta^0} < \theta < j \frac{2\pi}{N_\theta^0}$$

and will, in turn, process the first n_p pulses transmitted and received in this sector. If fewer than this number of pulses occurs, the data received from the sector in question is ignored.

If the radar operates with a pulse repetition frequency of f_r pulses per second, then to ensure that n_p pulses are transmitted and received in each elemental azimuth sector requires that:

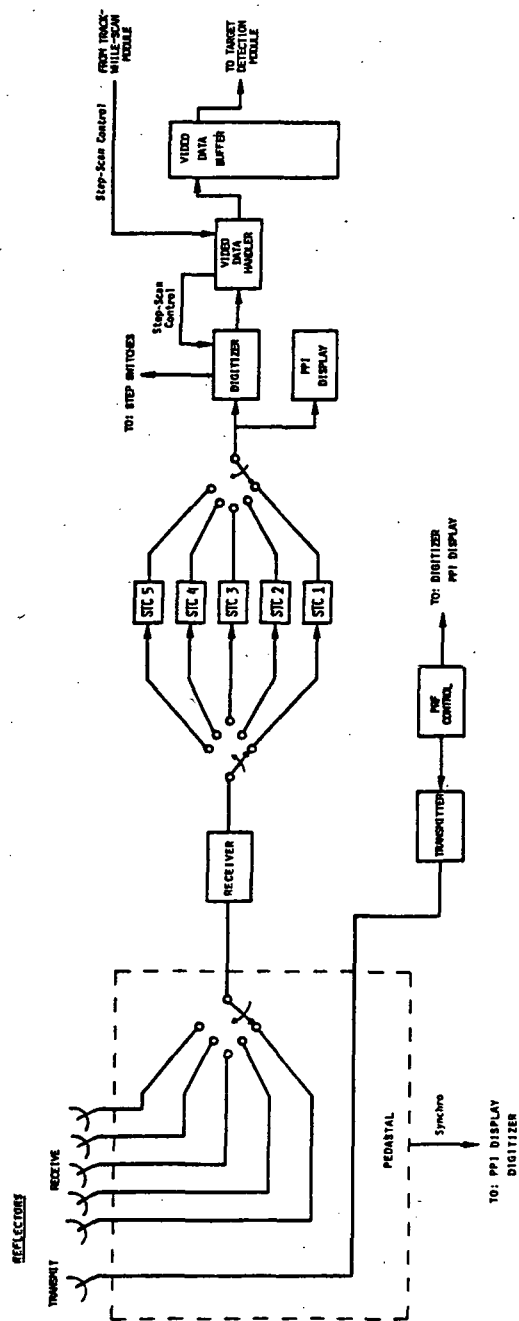


Figure 3-2. - Radar module.

$$\frac{n_p + 1}{f_r} < \frac{T_\sigma}{N_\theta^0},$$

or

$$f_r > \frac{N_\theta^0}{T_\sigma} (n_p + 1).$$

If $n_p = 8$, the current value of this parameter, then this constraint becomes

$$f_r > 1.15 \text{ kHz}.$$

On the other hand, in the interest of distributing pulses uniformly over each elemental sector and of minimizing the duty cycle of the radar, it is desirable to minimize f_r . The implication, therefore, is that f_r should, and is, chosen to be as small as possible, namely 1.15 kHz.

The radar video signal resulting from each of the n_p pulses processed in an elemental sector is sampled at a rate of τ^{-1} samples per second, where τ is the transmitted pulse width and can be either 0.5 μsec or 1.0 μsec . The sampling process commences T_b seconds after pulse transmission and continues until $N_r^0 = 64$ samples have been accumulated. Thus sampling occurs at the times

$$t(i) = T_b + (i-1) \tau, \quad i=1, \dots, N_r^0.$$

In geometric terms, the sampling process leads to radar returns from the set of ranges

$$r(i) = \frac{c}{2} t(i) = \frac{c}{2} T_b + (i-1) \frac{c}{2} \tau, \quad i=1, \dots, N_r^0.$$

Avoidance of ambiguous range problems requires, of course,

$$\frac{1}{f_r} > T_b + (N_r^0 - 1) \tau,$$

a constraint which is of little consequence for the short-range APAS radar.

As a result of the overall data acquisition process, one obtains a set of n_p radar samples, denoted $s_{i,j,k}(n)$, every $N_\phi T_\sigma$ seconds* which measure the reflected energy

*In the absence of beam-holding.

received from a region about the point whose coordinates are:

$$r(i) = \frac{c}{2} T_b + (i-1) \frac{c}{2} \tau$$

$$\theta(j) = (j-\frac{1}{2}) \frac{2\pi}{N_\theta^0}$$

$$\phi(k) = \phi_0(k),$$

where $\phi_0(k)$ is the nominal boresight angle of reflector k . These samples are immediately summed by the digitizer (pulse integration) to form the fundamental data of the radar module:

$$s_{i,j,k} = \sum_{n=1}^{n_p} s_{i,j,k}(n).$$

Pulse integration, of course, boosts the signal-to-noise ratio associated with the data upon which target detection decisions will be made. Since the integration process is noncoherent, the signal-to-noise ratio increase is somewhat less than 9 dB with $n_p = 8$.

3.1.1.2 Signal processing. The radar module has several modes of operation in which additional processing on the integrated samples $s_{i,j,k}$ may or may not be performed before these data are released to the target detection module. In the simplest mode, the so-called dual-mode, no further processing occurs. Thus the output of the module is:

$$s^0(i,j,k) = s_{i,j,k}$$

and the number of azimuth resolution sectors is

$$N_\theta = N_\theta^0 = 256.$$

In the second mode of operation, single-mode, the $s_{i,j,k}$ are combined pairwise to halve the number of resolution azimuth sectors to

$$N_\theta = \frac{N_\theta^0}{2} = 128.$$

The actual combining occurs in either of two ways: summing, in which

$$s^0(i,j,k) = s_{i,2j-1,k} + s_{i,2j,k};$$

and maximizing, in which

$$s^0(i,j,k) = \max \{s_{i,2j-1,k}, s_{i,2j,k}\}.$$

In general, superior performance is obtained with the latter option in single-mode operation.

3.1.1.3 Resolution and coverage. From the manner in which the $s^0(i,j,k)$ are formed, it is clear that the resolution achieved by the tracking data unit cannot exceed the limits imposed by the data acquisition and subsequent signal processing operations. Thus one cannot expect to resolve targets spaced closer than $c\tau/2$ meters in range or $2\pi/N_\theta$ radians in azimuth. In terms of actual parameters, the ultimate horizontal resolution of the system is given by the following table:

$N_\theta \backslash \tau$	0.5 μsec	1.0 μsec
128	75 m x 2.8°	150 m x 2.8°
256	75 m x 1.4°	150 m x 1.4°

In point of fact, the detection algorithms employed in the target detection module tend to decrease resolution, so that these numbers are truly lower bounds.

In the absence of transmitted power limitations, the maximum range of the system is constrained by the number N_r^0 , of video samples obtained, that is,

$$R_{\max} \approx \frac{c}{2} T_b + (N_r^0 - 1) \frac{c\tau}{2} + \frac{c\tau}{4}.$$

Its minimum range, of course, depends on the blanking time:

$$R_{\min} \approx \frac{c}{2} T_b - \frac{c\tau}{4}.$$

Thus one has, if

$$T_b = 2.5 \tau,$$

the following table:

τ	0.5 μ sec	1.0 μ sec
R_{\min}	150 m (0.08 nmi)	300 m (0.16 nmi)
R_{\max}	4930 m (2.66 nmi)	9860 m (5.33 nmi)

With reference to this and the preceding table, it has been found satisfactory to operate APAS with $\tau = 1.0 \mu$ sec and $N_0 = 128$. Although this particular combination contains the poorest overall resolution, it nevertheless provides the longer range and the lesser data rate. Target detection being a demanding task for the minicomputer, the latter factor, a low data rate, becomes quite significant in achieving real-time operation with a relatively low-cost, general-purpose computer.

Coverage in elevation, of course, is a function of the array of receiving antennas one employs. The current system configuration, described earlier in Section 2.1.1, is characterized as follows:

REFLECTOR k	BORESIGHT $\phi_0(k)$	VERTICAL BEAMWIDTH $\phi_\Delta(k)$
1	8°	4°
2	30°	18°
3	12°	7°

Although the upper limit of coverage is seen to be approximately 48° , the enhanced radar cross-section presented by high-elevation aircraft effectively moves this upper limit considerably higher. Although resolution in elevation is seen to be very coarse compared to horizontal resolution, it is nevertheless sufficient, when augmented by intelligent tracking, to provide the height-finding capability required of APAS.

3.1.1.4 Data handling. Output data from the digitizer are fed to the minicomputer in blocks of 64 16-bit words (via a standard DMA transfer), a single block being released following the interrogation of each of the N azimuth sectors. In any given block, the samples $s^0(i,j,k)$, $1 = 1, \dots, N_r^0$, is accorded 12 bits in word i . The remaining 4 bits per word, or 256 bits in all, are used to transfer supplementary information, namely azimuth sector j ; elevation beam k , and various status and error flags. The video data handler in Figure 3-2 strips off and interprets the supplementary bits and, according to their dictates, stores the video samples in a buffer memory.

According to the various parameters characterizing the radar module, a total of:

$$N_{\theta} N_r^0 N = 49,152$$

video samples are obtained per survey in dual-mode, 24,576 samples in single-mode. It will be seen that these samples are not accorded the same processing by the target detection module, some being ignored completely, some given only cursory treatment, and some afforded elaborate inspection. The primary purpose of the buffer memory into which the video samples are placed by the radar module is to allow the target detection module to proceed at an average pace, nominally:

$$\frac{N_{\theta} N_r^0 N}{T_{\sigma}} = 24,576 \text{ or } = 12,288$$

samples per second, instead of the pace determined by the most extensively processed samples.

Since the entire activity of radar interrogation proceeds sequentially, it is, in principle, unnecessary to require the digitizer to transfer the supplementary information described above. Doing so, however, considerably strengthens the digitizer/minicomputer interface, allowing the system, for example, to be self-synchronizing in the case of missing data. In addition, and ultimately of greater importance, is the monitoring capability provided by the error flags. By means of these flags, the minicomputer is made aware of undesirable perturbations to the system, the most common of which is wind loading the antenna assembly causing its pointing direction to deviate from nominal. Since this sort of error can seriously compromise fixed-window processing, data obtained under this condition are flagged by the digitizer and eliminated by the video data handler from further consideration.

3.1.2 Target Detection Module

The function of the target detection module is to identify returns in the radar video signal caused by aircraft while suppressing false returns caused by clutter, jamming,* and noise. The underlying philosophy upon which design of this module is based is simply stated: false targets have a greater potential for degrading APAS

*In this context, jamming refers to all electromagnetic interference, whatever the source.

performance then occasional missed targets. This philosophy, borne out in system testing, manifests itself in several ways in the discussion which follows.

The process of automatic target detection is accomplished in two sequential stages in APAS: hit report generation followed by target report generation. The first stage is concerned with finding the customary peaks in the received video signal; it is here that one attempts to eliminate the effects of clutter, jamming, and noise. The peaks, or hit reports, are then subjected to a second stage of processing in which they are clustered into target reports, the object being to account for the possibility that a single target can create sizeable returns in adjacent resolution cells.

The output of the target detection module are these target reports, described in terms of their location and magnitude, which are ultimately combined by the track-while-scan module into target tracks.

In what follows, the term resolution cell will refer to a fixed region of the space under APAS radar surveillance formed by

$$r(i) - \frac{c\tau}{4} < r < r(i) + \frac{c\tau}{4} ;$$

an azimuth sector:

$$\theta(j) - \frac{\pi}{N_{\theta}} < \theta < \theta(j) + \frac{\pi}{N_{\theta}} ;$$

and an elevation beam:

$$\phi_0(k) - \frac{\phi_{\Delta}(k)}{2} < \phi < \phi_0(k) + \frac{\phi_{\Delta}(k)}{2} ;$$

where the various quantities are defined in the preceding section. Resolution cells specify concretely the regions from which the video samples, $s^0(i,j,k)$, can be said to emanate and are a natural construct of the data acquisition procedure used in the radar module.

3.1.2.1 Clutter mapping. The underlying basis for hit-report generation in the target detection module is clutter mapping, wherein each successive video sample from a given resolution cell is compared to past sample values from that cell to eliminate any long-term bias. This is a standard radar technique which attempts to suppress the effects of close-in ground clutter, stationary objects, etc., in target detection. Briefly stated, the clutter-map indicates an historical average of the video samples received from each of the resolution cells, and it is the extent to which the latest video sample from a cell exceeds the clutter-map value for that cell that determines whether a hit report will be declared.

A clutter map consists, therefore, of a set of average values $\bar{s}(i,j,k)$, maintained for each of the resolution cells. When a new video sample, $s^0(i,j,k)$, becomes available for cell (i,j,k) , it is immediately reduced by the clutter-map value for that cell preparatory to any thresholding operation; one forms the reduced sample

$$\tilde{s}^0(i,j,k) = s^0(i,j,k) - \bar{s}(i,j,k).$$

Following thresholding, the clutter map value $\bar{s}(i,j,k)$ is updated according to the relation

$$\bar{s}(i,j,k) \leftarrow \bar{s}(i,j,k) + \tilde{s}^0(i,j,k)/8.$$

By means of this moving-average computation, the clutter map is made to adapt to a changing radar clutter environment.

At the time of system start-up, the clutter map is initialized at a value slightly above the maximum value of a video sample; i.e., at start-up,

$$\bar{s}(i,j,k) \leftarrow S_{\max} + 1 = 2041.$$

This approach prevents the flood of clutter hits that would otherwise occur were the map initialized at, say, zero. According to the updating relation above, the clutter map is allowed, rather, to decrease gracefully to steady-state values in the process preventing hit reports.

It is recalled from the previous section that there exist 24,576 resolution cells in the preferred APAS configuration. In principle, therefore, a clutter map would require a like number of memory locations for its storage in the minicomputer. Fortunately, however, in view of this memory demand, it is not necessary to clutter-map all cells, but only those which demonstrate returns consistently above the noise threshold of the radar.

To take advantage of this situation and reduce the memory requirements, the target detection module only clutter-maps patches of cells. As its name implies, a patch consists of a rectangular array of cells within a given elevation beam defined by minimum and maximum values of both the range bin index i and the azimuth sector index j . For a single beam therefore, a clutter map specification might take the form displayed graphically in Figure 3-3. Here there are only three patches shown for simplicity: one in the form of a ring surrounding the radar and intended to combat close-in ground returns; and the other two somewhat isolated and ostensibly covering stationary reflectors. By definition, all resolution cells outside the clutter-map specification carry the value zero:




Figure 3-3. - Clutter map.

$$\bar{s}(i,j,k) = 0 .$$

Although restricting a clutter map to patches within the various beams compromises somewhat the adaptability of APAS and renders it more airport-dependent, the reduced memory requirements allow the target detection module to reside in a general-purpose mini-computer. In the present APAS configuration, approximately 11,000 words of memory are available for clutter-mapping, a value which has proven satisfactory in practice.

3.1.2.2 Thresholding. The principal step in hit-report generation is the usual detection operation in which the reduced video sample is compared against a threshold to qualify it as indicating a target. In this regard, these distinct thresholds are employed by the target detection module: a clutter-map threshold, a mean-level threshold, and a noise threshold. Generally speaking, the first two are adaptive in nature while the last is fixed. (See, however, Section 3.1.2.3 below.)

The clutter-map thresholds are obtained directly from the clutter map. For a given cell (i,j,k) this threshold is given by:

$$\lambda_c(i,j,k) = 4 \bar{s}(i,j,k) .$$

Although a multiplier of four may seem extreme here, the time variability displayed by some types of clutter precludes the use of any lower value. One unfortunate consequence of having to use such a large multiplier is clear: whenever

$$\bar{s}(i,j,k) > \frac{2040}{5} ,$$

it will be impossible to declare a hit in cell (i,j,k) . One of the primary functions of the antenna screen described earlier is to minimize the number of cells in which this condition occurs.

The mean-level thresholds used in APAS constitute simply one variant of this well-known technique for combatting extended weather clutter and jamming. The general approach is to compare a video sample against those from adjacent cells obtained in the same scan; it is an analog of clutter-mapping in the spatial, as opposed to temporal, sense. The mean-level threshold for resolution cell (i,j,k) is computed from the video samples of cells about it, specifically cells $(i-2,j,k)$, $(i-1,j,k)$, $(i+1,j,k)$ and $(i+2,j,k)$.* Thus

*The obvious adjustments are made for end cells.

$$\lambda_m(i,j,k) = 4 \max \left\{ \max \{s^0(i-2,j,k), s^0(i+2,j,k)\}, \min \{s^0(i-1,j,k), s^0(i+1,j,k)\} \right\}.$$

Whereas the standard mean-level approach would typically use the average of the four indicated samples to compute a threshold, the higher threshold used in APAS was found necessary to combat jamming. (Although a recurring problem from, e.g., weather radars, jamming is nevertheless seldom so severe as to justify off-hand a system shutdown.) The APAS threshold avoids marginal shutdowns without seriously affecting detection performance.

Note in the threshold computation that the larger of the two video samples from cells adjacent to the target cell is effectively eliminated from consideration. This approach allows targets straddling two range bins to be detected; such targets are missed if a simple maximum of the four samples is used in computing the threshold. Note also that a multiplier of four is employed in computing λ_m as in the case of the clutter-map threshold λ_c . Experience indicates this to be a reasonable value.

The final type of threshold used in hit-report generation, namely noise thresholds, is primarily of importance in cells which are not clutter-mapped. For each range bin, i , of each elevation beam, k , a noise threshold, $\lambda_n(i,k)$, is set which reflects received noise level as modified by the sensitivity time control (STC) for that beam. Since there is an individual STC control for each beam, its attenuation characteristics tailored specifically to that beam, the noise thresholds are seen to depend potentially on both i and k as indicated. (These thresholds, incidentally, are empirically determined.)

With respect to the three thresholds defined above, then, a hit report is declared in cell (i,j,k) if and only if:

$$\tilde{s}^0(i,j,k) \geq \max \{ \lambda_c(i,j,k), \lambda_m(i,j,k), \lambda_n(i,k) \}.$$

This detection criterion has proven satisfactory in suppressing noise, most ground clutter, some weather clutter, and mild jamming. On the basis of APAS testing to date, it is deemed the most stringent condition one can impose without seriously degrading target detectability.

3.1.2.3 Hit-report suppression. Although the thresholding condition used by the target detection module constitutes a stringent test for generating a hit-report, it nevertheless proves inadequate in preventing the occurrence, from time to time, of excessively many hits caused, for example, by rain cells or severe jamming. Rather than

simply shut APAS down in these circumstances, two additional operations are performed as part of hit-report generation, namely threshold adjustment and blanking, to allow an orderly degradation in system performance followed by an automatic recovery. Roughly speaking, threshold adjustment suppresses marginal hits by raising the effective noise threshold. Until the mechanism of threshold adjustment is able to control the number of hits to a reasonable number, blanking is invoked to prevent an overload of hit reports on the track-while-scan module.

Both threshold adjustment and blanking (as well as track-while-scan processing) are based on a grouping of the $N_\theta = 128$ azimuth sectors into eight processing sectors, or simply octants. These octants, together with the $N_\phi = 3$ beams, coarsely subdivide the radar coverage volume into 24 regions, in each of which is monitored the number of hit reports occurring. Also associated with each of the regions is a threshold adjustment term, denoted $\Delta\lambda(\ell, k)$ and a blanking indicator $I(\ell, k)$. In operation, the noise threshold $\lambda_n(i, k)$ in the fundamental detection criterion is replaced by an effective threshold,

$$\lambda_n(i, k) + \Delta\lambda(\ell(j), k) ,$$

for resolution cell (i, j, k) ; here $\ell(j)$ denotes the octant containing sector j . Video samples satisfying this modified detection criterion continue to be declared hit reports; they are not, however, passed to the target report generation activity unless the corresponding blanking indicator is in a reset state.

Both the threshold adjustment term, $\Delta\lambda(\ell, k)$, and the blanking indicator, $I(\ell, k)$, are made to depend strictly on the number of hit reports occurring in octant ℓ and/or beam k . Depending on the precise pattern in which the hits occur (within the normal scanning pattern of the system), however, different actions are performed on these parameters. In the first instance, attention is focused on the octants. Initialized at zero, the threshold adjustment terms for the ℓ^{th} octant of all N_ϕ beams are incremented by 40 units (relative to a maximum video sample value of 2040) whenever the number of hits recorded in that particular octant of any single beam reaches a multiple of four. In this way, a large number of hits in this octant of any beam makes it more difficult for additional hits to occur in the same octant, whatever the beam.

Simultaneously, these same N_ϕ threshold terms are incremented by an additional 40 units whenever the total number of hits in the octant (summed over all beams) reaches a multiple of twelve. In this case, however, the hit counter is reinitialized to zero every third survey. As a result, this supplementary upward adjustment mechanism responds to a mild, but nevertheless significant, buildup of hit reports undetected by the first hit monitor.

Given this basic upward mobility of thresholds, a mechanism for downward adjustment is provided simply by decrementing all threshold adjustment terms by 10 units at the end of every third survey. The net result is a dynamic movement of thresholds up and down: a large number of hits causes a rapid movement upward; when the number of hits stabilizes, a slower downward movement ensues.

During the transition period in which the threshold adjustment terms for a given octant move upward in response to precipitation or jamming, the blanking indicators $I(\ell, k)$ for that octant (and all N_ϕ beams) may be set as a protective mechanism. Thus, whenever the cumulative number of hits recorded in any given octant during a single survey reaches eight, these indicators are set, thereby preventing any further hit reports in that octant from generating target reports. This status is maintained as long as the total number of hits in the octant in each successive survey exceeds eight. When this total eventually drops below eight, the blanking indicators are reset.

The measures for hit-report suppression just described apply on an octant basis. A similar threshold adjustment scheme, but applied on a beam basis, has not been found necessary. Beam-blanking, however, does prove beneficial, particularly during system startup while the clutter map is stabilizing. The general mechanism is the same: whenever the number of hit reports observed to occur during a given scan is found to be excessive, the corresponding beam is blanked, i.e., the $I(\ell, k)$ are set for all octants of that beam. In this case, however, the hit threshold is set equal to the number of target tracks currently maintained by the track-while-scan module plus eight; thus the threshold moves according to target density. Furthermore, beam-blanking, once invoked, requires five successive surveys in which the hit threshold is not exceeded before it is cleared.

The various techniques and parameter values introduced above are the result primarily of trial-and-error testing. Although they appear suited towards allowing a graceful system degradation in the case, e.g., of a thunderstorm, followed by system recovery, no claim is made that all are absolutely necessary not that the parameter values are optimal. It is asserted, however, that some such measures are required and that any working system must contain mechanisms similar to those described to qualify as being truly automatic.

3.1.2.4 Target-report generation. As a result of hit-report generation, each resolution cell (i, j, k) acquires a number, namely:

$$\max \{ 0, \tilde{s}^0(i, j, k) - \lambda_n(i, k) - \Delta\lambda(\ell(j), k) \},$$

which, if nonzero, signifies a hit report in that cell for the particular radar survey in question. The indicated magnitude of a hit report is the degree to which the reduced signal $\tilde{s}^0(i, j, k)$ exceeds the adjusted threshold $\lambda_n(i, k) + \Delta\lambda(\ell(j), k)$. Ideally, each

scan will see a target of interest produce at most one hit report in the appropriate resolution cell; and conversely, any hit report signifies a target of interest in its resolution cell. Even if it be assumed, however, that all hit reports are created by valid targets, the ideal situation just described does not obtain, because of range-straddle and azimuth-straddle, targets lying near the boundaries of resolution cells often produce hit reports in adjacent cells. It is the function of target report generation to deal with this phenomena. As a result, hit reports come to be clustered together into so-called target reports.

Conceptually, this clustering process involves two distinct steps: (1) determining the number of clusters (i.e., the number of target reports); and (2) specifying the hit reports comprising each cluster. For reasons of economy of storage and achievement of real-time operation, the target detection module performs these steps according to a very simple algorithm. Proceeding on a per scan basis, the module declares a target report each time a given hit report is a local maximum, i.e., its magnitude exceeds that of any other hit report in one of the eight resolution cells nearest to the cell in question (Figure 3-4). (As a result, for example, any isolated hit report produces a target report.) The cluster of hit reports comprising a target report, then, is simply the set of all hit reports, at most nine, which enter into the identification of the existence of that target report.

From the description just given, it is clear that target report generation, viewed as an intermediate step between the generic activities of detection and tracking, serves to reduce the number of objects, real or otherwise, which will be tracked. In this sense, and in this sense only, can the process be considered to perform clutter suppression. Furthermore, there is inherent in any clustering algorithm a loss of radar resolution; thus, for example, the particular clustering technique just described will fail to distinguish targets occupying resolution cell (i,j,k) and $(i+1,j,k)$. Such loss of resolution is justified in APAS, however, by the dictum mentioned earlier: false targets (or ghost targets) are ultimately more damaging to system performance than missed targets.




Figure 3-4. - Array of cells used to generate target reports.

3.1.3 Track-While-Scan Module

The final module of the surveillance radar unit is responsible for converting target-reports into tracks, an operation typically termed track-while-scan (TWS) processing to reflect the dual functions of surveillance and tracking. Conceptually, TWS processing is straightforward: one simply assembles target reports into tracks according to their positional relationship to one another. Unfortunately, a number of factors greatly complicate this process, some generic to all TWS processors, some peculiar to APAS. These complications are addressed at the appropriate time in the sections which follow.

One problem, however, which should be mentioned at the outset is analogous to that mentioned earlier in the context of target-report generation, namely the constraint of real-time operation. So that traffic advisories might be delivered in a timely fashion, it is necessary for the TWS processor to process target reports as they are received. Neither is there a look-ahead possibility (i.e., postponed decision making) nor any realistic opportunity for retroactive processing. It is within this fundamental constraint that the TWS module must operate.

Given that target reports are to be processed as they are received, TWS processing divides conceptually into two basic phases: an association phase in which target reports are allocated to existing tracks and an initiation/updating phase in which these target reports are actually incorporated into tracks. Complementing this two-phase, largely mechanistic procedure are a series of functions performed by the TWS module of a more judgmental nature which serve to manage the evolution of tracks, evaluate their goodness (or firmness), and decide on their fitness to appear in a traffic advisory. The basic processing structure which allows these various tasks to occur in an orderly fashion is provided by the processing sectors or octants.

Thus, batches of target reports from a given processing sector are examined as a group with reference to the tracks which lie in or near that sector. When each such batch of target reports has been processed, the tracks in question are examined retrospectively and evaluated according to whether or not they have been updated. Those which have may become, or remain candidates for reporting in traffic advisories; those which have not may enter a coasting state and, when this occurs, either lose, or fail to gain reporting status.

The following subsections serve to document these activities more fully. The sections are arranged according to the basic functions performed by the TWS module: tracking, association, classification, management, and evaluation. As implied above, tracking and association, the two basic activities of TWS processing, are performed no differently in APAS than in any other TWS radar system. The key to the successful tracking of uncooperative targets in APAS lies in exploiting the relative orderliness of

an aircraft traffic flow, even in the relatively unconstrained environment of an uncontrolled airport. This orderliness allows tracks to be classified in such a way that some inference can be made of their future behavior. When tracks fail to behave properly, they can be denied reporting status or, in the more severe cases, terminated. The net result is a further reduction in false target reports over that achieved by the target detection module.

In closing this introduction, it is useful to note that the output of the TWS module, i.e., tracks, actually consists of a series of descriptive data blocks, each such block constituting a track. Information contained in the individual data blocks pertains to the following items:

- target position and velocity
- target steering controls (with reference to maneuvers)
- target radar-return strength
- target classification (or pattern state)
- target goodness indicators
- track reporting status.

These data blocks are maintained in memory in the APAS minicomputer and are periodically interrogated by the voice response unit for information to include in traffic advisories.

3.1.3.1 Tracking. The most convenient of the functions of the TWS module to describe first is tracking, the process by which target reports are used either to update existing tracks or to initiate new tracks. The tracking algorithm used in APAS is a version of what is commonly called an $\alpha - \beta$ tracker, the sort of algorithm commonly found employed in similar aircraft surveillance systems. Although lacking the sophistication of Kalman filters, the $\alpha - \beta$ tracker also avoids some of the problems of the former class of filters, e.g., numerical stability, and requires far fewer computations in its operation. Furthermore, experience has proven that the primary benefit to be provided by any tracker in the APAS application is the extent to which it improves the association process, that is, the ability of the TWS module to maintain track on a target accurately in spite of target maneuvers and misleading target reports. So important is this characteristic that the performance of the tracker in providing smooth estimates of target position and velocity, the usual criterion for judging such filters, is a secondary consideration by comparison.

The discussion which follows develops the particular version of the $\alpha - \beta$ tracking filter used in APAS according to the standard state-space approach. Beginning with a general kinematical model for aircraft motion, the development results in the customary two sets of filter equations, one for predicting target state on the basis of current estimated state and one for estimating target state on the basis of current predicted

state and a new target report. In contrast to the usual $\alpha - \beta$ tracker, that described here also requires an underlying statistical model for the purpose of association, the topic of the next section. It is also noted that the present tracking filter incorporates rudimentary adaptivity features in which the filter parameters are allowed to vary and in which target maneuvers are permitted in predicting target state.

To begin, target (aircraft) motion is modelled by the following set of equations:

$$x(t + \tau) = x(t) + \tau \dot{x}(t) + \frac{\tau^2}{2} \ddot{x}(t)$$

$$\dot{x}(t + \tau) = \dot{x}(t) + \tau \ddot{x}(t)$$

$$y(t + \tau) = y(t) + \tau \dot{y}(t) + \frac{\tau^2}{2} \ddot{y}(t)$$

$$\dot{y}(t + \tau) = \dot{y}(t) + \tau \ddot{y}(t)$$

$$z(t + \tau) = z(t) + \tau \dot{z}(t),$$

where the overdots signify this differentiation.

With (x, y, z) denoting target position (with respect to the radar), these equations relate the position and velocity of a target at time $t + \tau$ to its position, velocity, and acceleration at time t . The model, of course, is simply a series expansion truncated at second-order terms with respect to the horizontal coordinates x and y and at first-order terms with respect to the vertical coordinate z .

A useful refinement to the model involves eliminating the horizontal acceleration variables \ddot{x} and \ddot{y} in favor of a target-oriented set of coordinates which are more susceptible to precise characterization. Thus, if v denotes target velocity,

$$v^2 = \dot{x}^2 + \dot{y}^2 + \dot{z}^2$$

and ψ denotes heading,

$$\psi = \tan^{-1} \frac{\dot{y}}{\dot{x}}$$

then the required transformation is

$$\ddot{x} = \frac{v \cos(\psi)}{s} \dot{v} - s \sin(\psi) \dot{\psi}$$

$$\ddot{y} = \frac{v \sin(\psi)}{s} \dot{v} + s \cos(\psi) \dot{\psi},$$

where

$$s^2 = \dot{x}^2 + \dot{y}^2 = v^2 - \dot{z}^2.$$

Here in-track acceleration,

$$a \triangleq \dot{v}$$

and turn rate,

$$w = \dot{\psi},$$

become control variables in the model together with rate of climb,

$$p \triangleq \dot{z}.$$

If one adopts the approximations, generally valid for general aviation aircraft, that vertical velocity is negligible compared to horizontal velocity, i.e., $s \approx v$, and that in-track acceleration is negligible, i.e., $a \approx 0$, then the final target dynamical model is obtained:

$$x(t + \tau) = x(t) + \tau \dot{x}(t) - \frac{\tau^2}{2} \dot{y}(t) \omega(t)$$

$$\dot{x}(t + \tau) = \dot{x}(t) - \tau \dot{y}(t) \omega(t)$$

$$y(t + \tau) = y(t) + \tau \dot{y}(t) + \frac{\tau^2}{2} \dot{x}(t) \omega(t)$$

$$\dot{y}(t + \tau) = \dot{y}(t) + \tau \dot{x}(t) \omega(t)$$

$$z(t + \tau) = z(t) + \tau \rho(t).$$

The model is seen to involve five state variables: x , \dot{x} , y , \dot{y} , and z ; and two control variables ω and ρ .

In what follows, it is useful to express the model equations in matrix form; to this end, if

$$\underline{\bar{X}} \triangleq \begin{bmatrix} x \\ \dot{x} \\ y \\ \dot{y} \\ z \end{bmatrix}$$

represents target state and

$$v = \begin{bmatrix} \omega \\ \rho \end{bmatrix}$$

target control, then

$$\underline{\bar{X}}(t + \tau) = A_\tau + B_\tau(\underline{\bar{X}}(t)) v(t),$$

where the matrices A_τ and $B_\tau(x)$ have the obvious definitions.

Given the target dynamical model, attention is now shifted to a complementary observation model. From the preceding section, it is recalled that a target report consists of a cell identification (i,j,k) and the maximum of nine hit reports associated with that cell and its eight nearest neighbors. This target report gives rise to an observation simply by computing the centroid of the hit reports (imagined as point masses located at the centers of their respective resolution cells). One thus arrives at a nominal position (r^0, θ^0, ϕ^0) for this observation.* In addition, there is a time t^0 associated with the observation,

$$t^0 = \frac{T_\sigma}{2\pi} \theta^0.$$

The spherical coordinates (r^0, θ^0, ϕ^0) are converted to rectangular coordinates (x^0, y^0, z^0) according to the usual transformation:

$$x^0 = r^0 \cos(\theta^0) \cos(\phi^0)$$

$$y^0 = r^0 \sin(\theta^0) \cos(\phi^0)$$

*It will be noted that, according to this definition, $\phi^0 = \phi_0(k)$, the boresight elevation of reflector k. For the moment this temporary definition is acceptable.

$$z^0 = r^0 \sin(\phi^0)$$

to agree with the coordinate frame of the target dynamical model. Again, for convenience, let the vector

$$Y^0 = \begin{bmatrix} x^0 \\ y^0 \\ z^0 \end{bmatrix}$$

denote the position of an observation.

To create the dynamical and observational models just described a tracking filter, imagine an existing track specified by a latest state estimate \bar{X}^* obtained at time t^* . If v^* denotes the target control predicted to be in effect at this time, then the model may be used to generate a predicted state $\bar{X}^*(t)$ at any sensible time $t > t^*$ according to

$$\bar{X}^*(t) = A_{t-t^*} \bar{X}^* + B_{t-t^*}(\bar{X}^*)v^*(t), \quad (P)$$

Let there be given, then, an observation Y^0 occurring at time t^0 which, it appears, is attributable to the target in question. The process of updating the track of this target involves an adjustment of its predicted state $\bar{X}^*(t^0)$ to arrive at a new estimated state:

$$\bar{X}^* = \bar{X}^*(t^0) + G_{t^0-t^*} (Y^0 - H\bar{X}^*(t^0)), \quad (E)$$

where H is the matrix

$$H = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

The matrix G_τ , the so-called gain matrix, is the distinguishing characteristic of an $\alpha - \beta$ tracker and is given by

$$G_\tau = \begin{bmatrix} \alpha & 0 & 0 \\ \beta/\tau & 0 & 0 \\ 0 & \alpha & 0 \\ 0 & \beta/\tau & 0 \\ 0 & 0 & \alpha' \end{bmatrix}.$$

The time, t^* , associated with this new estimate is chosen to synchronize the position (x^*, y^*) with the fundamental system clock, i.e., the rotating antenna. Thus, t^* is made to satisfy a relation of the form

$$\tan \left(2\pi \frac{t^*}{T_\sigma} \right) = \frac{y^*}{x^*}.$$

The parameters of the tracking filter-- α , β , and α' --must be chosen to provide the proper filter response. Generally speaking, making the filter parameters large increases responsiveness of the filter to new data and, potentially, allows it to cope with target maneuvers. At the same time, however, this approach allows any errors in the observed position (x^0, y^0, z^0) to be reflected in the estimated state of the target. It has been generally accepted that setting

$$\alpha = 0.75$$

and

$$\beta = \frac{2}{2 - \alpha} = 0.45$$

represents a good compromise in this regard with respect to horizontal target motion. The parameter α' , pertaining to vertical motion, is set to a value which prohibits radical changes in z :

$$\alpha' = 0.20.$$

This particular choice and, indeed, the overall treatment of the z -coordinate are motivated by the generally level flight displayed by aircraft and the poor ability of the system to measure elevation.

In operation, one initiates a track with an isolated radar observation (x^0, y^0, z^0) obtained at time t^0 :

$$x^* = x^0$$

$$\dot{x}^* = 0$$

$$y^* = y^0$$

$$\dot{y}^* = 0$$

(I)

$$z^* = z^0 ,$$

with

$$t^* = t^0 .$$

As succeeding observations which are judged pertinent to this track become available, the track is updated as indicated: first the predicted state of the track is computed according to (P) and then the new estimated state according to (E). In computing the predicted state of a track, the controls-- ω and ρ --are for the most part set to zero, signifying a general lack of knowledge as to what maneuvers might be expected of the target in question. In the case of tracks which appear to represent pattern aircraft, however, the control variables are often purposefully chosen to anticipate pattern behavior. The manner in which the TWS module exploits this possibility is discussed later in Section 3.1.3.4 .

To this point in the discussion, no statistical model has been mentioned, there being no explicit need for such in the definition of an $\alpha - \beta$ tracker. Such a model is important, however, for association, and before this topic is treated in the next section the errors associated with target state estimates and predictions must be characterized. To this end, it is first noted that the major sources of error in the $\alpha - \beta$ tracker are the following:

- range measurement;
- azimuth measurement;
- elevation measurement;
- initial velocity estimate;
- turn-rate prediction;
- rate-of-climb prediction.

For simplicity, the errors committed with respect to these variables are assumed to obey whatever statistical regularity and independence assumptions are required below and to be characterized by zero means and fixed variances.

Thus the components of the observation (r^0, θ^0, ϕ^0) associated with a given target report in resolution cell (i, j, k) possesses errors with variances given by

$$\sigma_{r0}^2 = \frac{1}{2\pi} \left(\frac{c\tau}{2} \right)^2 ;$$

$$\sigma_{\theta}^2 = \frac{1}{2\pi} \frac{2\pi}{N_{\theta}}^2 ;$$

$$\sigma_{\phi}^2 = \frac{1}{2\pi} \phi_{\Delta}(k)^2 .$$

If, then

$$Q \triangleq \text{diag } \sigma_r^2, \sigma_{\theta}^2, \sigma_{\phi}^2$$

and

$$C(r, \theta, \phi) = \begin{bmatrix} \cos(\theta) \cos(\phi) & -r \sin(\theta) \cos(\phi) & -r \cos(\theta) \sin(\phi) \\ \sin(\theta) \cos(\phi) & r \cos(\theta) \cos(\phi) & r \sin(\theta) \sin(\phi) \\ \sin(\phi) & 0 & r \cos(\theta) \end{bmatrix} ,$$

then standard linearization techniques show that the covariance matrix of the observation \bar{Y}^0 is given by

$$R(\bar{Y}^0) = C(r^0, \theta^0, \phi^0) Q C(r^0, \theta^0, \phi^0)^T .$$

With respect to the remaining sources of error, one assumes that errors committed by initializing \dot{x}^* and \dot{y}^* at zero have same variance,

$$\sigma_s^2 = \frac{1}{2\pi} (s_{\max} - s_{\min})^2 ,$$

where s_{\max} and s_{\min} are relatively unconstrained design parameters reflecting the maximum and minimum speeds, respectively, to be displayed by a general aviation aircraft. It then follows that the error covariance matrix of the first state estimate (cf. formula (I)) is given by

$$P^* = \begin{bmatrix} R_{11}(Y^0) & 0 & R_{12}(Y^0) & 0 & R_{13}(Y^0) \\ 0 & \sigma_s^2 & 0 & 0 & 0 \\ R_{21}(Y^0) & 0 & R_{22}(Y^0) & 0 & R_{23}(Y^0) \\ 0 & 0 & 0 & \sigma_s^2 & 0 \\ R_{31}(Y^0) & 0 & R_{32}(Y^0) & 0 & R_{33}(Y^0) \end{bmatrix}$$

where the R_{ij} denote the elements of the matrix R .

The error covariance matrices of subsequent state predictions and estimates then follow the customary two-step recursion relationship

$$P^*(t) = A_{t-t^*} P^* A_{t-t^*}^T + B_{t-t^*} (\bar{X}^*) S B_{t-t^*} (\bar{X}^*)^T$$

and

$$P^* = (I - G_{t0-t^*} H) P(t^0) (I - G_{t0-t^*} H)^T + (G_{t0-t^*} R/Y^0) G_{t0-t^*}^T.$$

Here

$$S = \text{diag}(\sigma_\omega^2, \sigma_p^2)$$

is a diagonal matrix whose elements are given by

$$\sigma_\omega^2 = \frac{1}{2\pi} \omega_{\text{MAX}}^2$$

and

$$\sigma_p^2 = \frac{1}{2\pi} \rho_{\text{max}}^2$$

analogous to σ_s^2 . By means of the above formulas, one has available at all times in the evaluation of a track some measure of the uncertainty associated with current state estimates and predictions.

While keeping the basic structure of the tracking filter described in the preceding paragraphs intact, several minor modifications are made in APAS to enhance tracking performance. The first concerns the transient response of the filter and its effect on track acquisition, i.e., the initial stages of establishing a track. If

$(x^0(1), y^0(1), z^0(1))$ and $(x^0(2), y^0(2), z^0(2))$ denote the first two observations forming a track, occurring at time $t^0(1)$ and $t^0(2)$, respectively, the formulas (I), (P), and (E) above imply that

$$x^*(2) = x^0(1) + \alpha(x^0(2) - x^0(1))$$

$$\dot{x}^*(2) = \beta \frac{x^0(2) - x^0(1)}{t^0(2) - t^0(1)}$$

$$y^*(2) = y^0(1) + \alpha(y^0(2) - y^0(1))$$

$$\dot{y}^*(2) = \beta \frac{y^0(2) - y^0(1)}{t^0(2) - t^0(1)} .$$

At this stage in establishing a track, however, it is rather pointless to attempt any filtering; more fruitful is to set $\alpha = 1$ and

$$\beta = \frac{\alpha^2}{2 - \alpha} = 1$$

to give

$$x^*(2) = x^0(2)$$

$$\dot{x}^*(2) = \frac{x^0(2) - x^0(1)}{t^0(2) - t^0(1)}$$

$$y^*(2) = y^0(2)$$

$$\dot{y}^*(2) = \frac{y^0(2) - y^0(1)}{t^0(2) - t^0(1)} .$$

Having then obtained decent initial estimates of both position and velocity, the filter parameters can henceforth revert at their previously stated values.

The second modification, somewhat more complex than the first, addresses the generally marginal vertical tracking performance of the original filter. Because an aircraft can often be detected in multiple elevation beams of the radar module, one finds that its track bounces excessively in spite of the narrow bandwidth filtering produced by setting $\alpha' = 0.20$. To treat this problem, one defines another observable, namely strength, which is defined as the sum of the hit reports comprising a target report. Denoted h^0 , this variable is related to the radar cross-section of a target and its position with respect to the antenna pattern of the beam in which the target is

detected. At the time of track initiation, one sets

$$h^* = h^0$$

as the estimated strength of the track. As later observations are associated with the track, predicted strength is given by

$$h^{\wedge}(t) = h^*, \quad t > t^*,$$

and a new estimated strength by

$$h^* = h^{\wedge}(t^0) + \alpha''(h^0 - h^{\wedge}(t)),$$

where $\alpha'' = 0.25$.

To relate, now, this new variable to height estimation, suppose that an observation (r^0, θ^0, ϕ^0) is associated with a given track. If $\phi^{\wedge}(t)$ is the predicted elevation of the track and $h^{\wedge}(t)$ its predicted strength, one replaces the observed elevation ϕ^0 by the new elevation

$$\phi^0 \leftarrow \frac{h^0 \phi^0 + h^{\wedge}(t^0) \phi^{\wedge}(t^0)}{h^0 + h^{\wedge}(t^0)},$$

a weighted average of the observed elevation, ϕ^0 , and the predicted elevation, $\phi^{\wedge}(t^0)$. Thereafter, filtering occurs as before. This straightforward device allows added weight to be given to observations occurring in the elevation beam most directly illuminating a target and correspondingly less weight to side observations.

It should be noted that the two modifications to the original $\alpha - \beta$ tracker affect inconsequentially the computation of error covariance matrices. The first is accounted for in the general formulas for these matrices given earlier and the second, if anything, reduces the variance term σ_{ϕ}^2 .

3.1.3.2 Association. The process of observation-to-track association in the TWS module attempts to reproduce the behavior of a human operator in tracking targets visually on a radar display. In this case, one mentally extrapolates a track ahead and then looks for radar returns occurring near this extrapolated position. A suitable return having been observed, the track is extended and the extrapolation process begun anew.

Stated in terms of the TWS module, the association process presents itself as follows. Let there be M tracks in existence, labelled $i=1, \dots, M$, and let there be

given a new set of N observations, labelled $j=1, \dots, N$.^{*} For each observation j , one computes the predicted position of each track, i.e., the vector $HX_i^{\wedge}(t_j^0)$. Comparing each observed position Y_j^0 with the corresponding predicted position $HX_i^{\wedge}(t_j^0)$ for track i allows one to select an observation to update the track. In the next general case, the net result of the association process is to dichotomize the set of observations into two classes; the first consisting of observations used to update existing tracks and the second of observations not so used. Observations of the latter type become, by default, the seeds of new tracks.

Several features of manual tracking are notable towards realizing automatic observation-to-track association:

- In the search for a new return from a target being tracked, most of the radar display is ignored; that is, radar returns far removed from the anticipated target position are dismissed off-hand.
- In the case of competing returns, the one closest to where the target is anticipated to be is chosen to extend its track.
- In the case of neighboring targets, returns are allocated to tracks in a way which appears most appropriate.

The approach to association described in the following paragraphs attempts to reflect these features.

The observations just made suggest the necessity of a metric to define the distance between an observation and a predicted target position. That is, one must measure the length of the vector

$$Y^0 - HX^{\wedge}(t^0)$$

appearing in formula (E) of the preceding section. For the moment, suppose that a length measure is defined, denoted by $\Delta(HX^{\wedge}(t^0), Y^0)$.

Given, then, the set of M existing tracks and the set of N new observations previously defined, one computes the quantities

$$\delta_{i,j} = \Delta(H\bar{X}_i^{\wedge}(t_j^0), Y_j^0) .$$

^{*}The set of observations here are to be thought of as emanating from a given processing sector.

For each track i , one next ranks the $\delta_{i,j}$ according to decreasing magnitude. This ranking leads to a tableau of the form

1	$j_1(1)$	$j_1(2)$. . .	$j_1(N)$
2	$j_2(1)$	$j_2(2)$. . .	$j_2(N)$
.	.	.		.
.	.	.		.
.	.	.		.
M	$j_M(1)$	$j_M(2)$. . .	$j_M(N)$

in which a given row, labelled by track number i , consists of observation numbers $j_i(1)$ arranged so that

$$i, j_i(\ell) \leq i, j_i(\ell+1).$$

Stated qualitatively, observation $j_i(\ell)$ is preferable to $j_i(\ell+1)$ for updating track i .

The next step is to truncate each row of the tableau to only those observations for which

$$\delta_{i,j} < \delta_{\max},$$

thereby disqualifying immediately from consideration for any track all observations remote from it. If this disqualification process should leave row i of the tableau empty, then track i will not be updated from the current set of observations.

Empty rows then having been removed (conceptually, anyway) from the tableau, one is left ideally with the situation in which the leading elements of the individual rows are all distinct. In this case, one simply uses observation $j_i(1)$ to update track i . Unfortunately, this situation does not occur if two or more tracks are competing for the same observation. To resolve such disputes, the following algorithm is employed:

- (1) Set $i \leftarrow 0$.
- (2) Set $i \leftarrow i + 1$.
- (3) If $i > M$, stop.

- (4) Set $i' \leftarrow i$.
- (5) Set $i' \leftarrow i' + 1$.
- (6) If $i' > M$, go to step (2).
- (7) If $j_{i'}(1) \neq j_i(1)$, go to step (5).
- (8) With

$$\delta_{ij} \triangleq \frac{\delta_{i,j}}{L_i},$$

where L_i is the length of track i , i.e., the number of times track i has been updated (including its initiation), compute and compare

$$\tilde{\delta}_{i,j_i(1)} \quad \text{and} \quad \tilde{\delta}_{i',j_{i'}(1)}.$$

If the former is smaller, delete $j_{i'}(1)$ from row i' and shift the row left; if not, delete $j_i(1)$ from row i and shift this row left.

- (9) Go to step (1).

The objective of this algorithm is to allow track i , instead of track i' , to seize observation j if either the observation is distinctly closer to the former track or this track has persisted longer than the latter. The point, of course, is to prevent a tentative track with only a few updates to persist at the expense of a well-established track.

To complete the specification of the association process, the metric used therein must be defined. To do so, reference is made to the statistical model of the preceding section in which it was stated (implicitly) that the covariance matrix of the (random) vector

$$Y^0 = \hat{H}X(t^0)$$

is given by

$$D(Y^0, t^0) = R(Y^0) + HP^{\wedge}(t^0)H^T.$$

Standard procedure suggests immediately that an appropriate metric is then given by

$$\Delta(H\bar{X}^{\wedge}(t^0), Y^0)^2 = (Y^0 - H\bar{X}^{\wedge}(t^0))^T D(Y^0, t^0)^{-1} (Y^0 - H\bar{X}^{\wedge}(t^0)).$$

Geometrically, the relation

$$\Delta(H\bar{X}^{\wedge}(t^0), Y) = \delta$$

defines a set of points Y which form an ellipsoid in three dimensional space centered at the point $H\bar{X}^{\wedge}(t^0)$. This fact is exploited further below.

As implemented in the TWS module, the association algorithm just described is modified somewhat to reduce storage and computation demands. In the first place, only two columns are retained in the tableau constructed above; that is, only the two best observations are considered for each track. The effect is to reduce the size of the tableau (i.e., the storage required for it) at no noticeable decrease in performance.

A second, more significant, modification concerns the predicted states $X^{\wedge}(t^0)$. Given the estimated state X^* of a track at time t^* , one can predict the series of times at which the boresight azimuth of the radar antenna structure will coincide with the target azimuth by solving the equation

$$\theta^{\wedge}(t) = \frac{2\pi}{T_{\sigma}} (t - t^*) \pmod{2\pi}$$

for its sequence of roots. If these solutions--predicted observation times--are denoted $t^{\wedge}(n)$, $n=1,2,3,\dots$, in increasing order, then one is led to believe that the state of the track when it is next observed will be given approximately by $X^{\wedge}(t^{\wedge}(n))$ if that observation occurs n scans after time t^* . Stated another way, observations received after time t^* will come to be used to update the track only if the observation time t^0 is near some predicted observation time $t^{\wedge}(n)$, in which case the predicted state $X^{\wedge}(t^0)$ will be near $X^{\wedge}(t^{\wedge}(n))$.

To exploit this observation, the rule is adopted that an observation is allowed to associate with a track if and only if: (1) the time of the observation t^0 is near some predicted observation time $t^{\wedge}(n)$ and (2) the position of the observation Y^0 is near the predicted position $H\bar{X}^{\wedge}(t^{\wedge}(n))$ (as opposed to the predicted position $H\bar{X}^{\wedge}(t^0)$). The benefits of this modified approach to association are twofold: it provides a convenient mechanism for monitoring the evaluation of a track, and it reduces the workload connected

with computation of the distances between tracks and observations.

The benefits appear as follows. Given that a track has just been updated (or initiated) at time t^* , the next predicted observation time, $t^{\wedge}(1)$, together with predicted position

$$Y^{\wedge}(1) = HX^{\wedge}(t^{\wedge}(1))$$

and the matrix

$$D(1)^{-1} \triangleq D(Y^{\wedge}(1), t^{\wedge}(1))^{-1},$$

and computed. Since the set of points Y for which

$$\Delta(Y^{\wedge}(1), Y)^2 = (Y - Y^{\wedge}(1))^T D(1)^{-1} (Y - Y^{\wedge}(1)) \leq \delta_{MAX}^2$$

is an ellipsoid about the point $Y^{\wedge}(1)$, one can compute the azimuthal angles subtended by this ellipse and therefrom two times

$$t^{\wedge}_{MIN}(1) < t^{\wedge}(1) < t^{\wedge}_{MAX}(1),$$

with the property that

$$\Delta(Y^{\wedge}(1), Y^0) \leq \delta_{MAX}$$

only if

$$t^{\wedge}_{MIN}(1) < t^0 < t^{\wedge}_{MAX}(1).$$

That is, the observation in question can associate with the given track only if t^0 lies on the indicated time interval. In effect, the time interval, when expressed geometrically in terms of angles, defines an association sector for the track such that observations from a processing sector which does not at least overlap the association sector cannot associate with the track.

The track in question just having been updated, therefore, and its association sector defined, the track enters a quiescent state until the current observation sector comes to overlap its association sector. When this occurs, the track enters an active state and remains so until it is either updated, requiring the computation of a new association sector, or bypassed, i.e., the current processing sector moves beyond the

existing association sector for the track. In the latter case, note is taken of the track having been bypassed and a new association sector constructed using the next predicted observation time, $t^{\wedge}(2)$. The process can repeat indefinitely until terminated by the management function (see Section 3.1.3.4).

Given that the track is an active state, then, one computes the $\Delta(Y^{\wedge}(1), Y_j^0)$ for the set of new observations and association proceeds as before. An additional benefit of the modified approach now becomes apparent, the possibility of delayed association. If $Y^{\wedge}(1)$ is situated such that the association sector for the track straddles two processing sectors, then it is possible that the track, even though it seizes an observation from in the first of the sectors, will find a better one in the next sector. In this situation, the track is not updated and the observation in question left in the pool of observations for reassessment with the observations from the next processing sector. In this way, tracks near the boundary of two processing sectors are not discriminated against in the association process.

The computational advantages of the new procedure lie in the elimination of repeated calculation of the matrix $D(Y^0, t^0)^{-1}$ for each observation-track pair. Now

$$D(n)^{-1} = D(Y^{\wedge}(n), t^{\wedge}(n))^{-1}$$

need be computed only once for each track each scan. Although it is now necessary to compute the $t^{\wedge}(n)$, $t_{\text{MIN}}^{\wedge}(n)$, and $t_{\text{MAX}}^{\wedge}(n)$, this is straightforward. Indeed, the first sequence of times is given approximately as follows: If

$$\Delta T^{\wedge}(n) \triangleq \frac{T_{\sigma}}{2\pi} (\theta^{\wedge}(n T_{\sigma}) - \theta^*),$$

then

$$t^{\wedge}(n) \approx t^* + (2\pi n + \theta^{\wedge}(n T_{\sigma}) + \Delta t^{\wedge}(n)) - \theta^* \frac{T_{\sigma}}{2\pi}.$$

Derivation of formulas for the times $t_{\text{MIN}}^{\wedge}(n)$ and $t_{\text{MAX}}^{\wedge}(n)$ is a laborious exercise in analytic geometry (and is not reproduced here). It should be noted, however, that if

$$\Delta(Y^{\wedge}(n), 0) < \delta_{\text{MAX}},$$

i.e., the association ellipsoid contains the origin, then by definition

$$t_{\text{MIN}}^{\wedge}(n) = t^{\wedge}(n) - T_{\sigma}/2 \text{ and } t_{\text{MAX}}^{\wedge}(n) = t^{\wedge}(n) + T_{\sigma}/2.$$

The final comment to be made concerning the association process involves the measured elevation of an observation, ϕ_j^0 . It is recalled that this measurement is modified before the observation in question is used to update a track; it is likewise modified prior to attempting to associate it with a given track. That is, in computing the position of observation j on rectangular coordinates, and thence the distance $\delta_{i,j}$ for track i , one replaces, temporarily, the measured elevation, ϕ_j^0 , by the current predicted elevation, $\phi_i^{\wedge}(t_i(n))$. The net effect is to remove elevation as a determining factor in association as described above, a useful modification in view of the general uncertainty in elevation measurement and the tendency for targets to appear in multiple elevation beams.

Lest this modification lead to false association, however, a preliminary check is performed to ensure that ϕ_j^0 and $\phi_i^{\wedge}(t_i(n))$ are in the same elevation beam. Thus, if ϕ_j^0 is detected in beam k , it is required, that

$$|\phi_j^0 - \phi_i^{\wedge}(t_i(n))| \leq \phi_{\Delta}(k) \left(1 + \left(\frac{3200}{r_j^0}\right)^4\right),$$

where $\phi_{\Delta}(k)$ is the vertical beam width of beam k . If this condition is not met, one sets

$$\delta_{i,j} = \delta_{\max},$$

ensuring that observation j will be deleted from row i of the association tableau.

3.1.3.3 Classification. With minor additions, primarily having to do with track termination and steering, the TWS algorithms presented to this point allow the construction of a complete, general-purpose TWS processor which will provide satisfactory performance in a relatively clean radar environment. By satisfactory performance, it is emphasized, is meant the absence in traffic advisories of reports on nonexistent aircraft.

Unfortunately, such radar environments are apparently the exception rather than the rule; severe problems were encountered at each of the three airports at which the experimental APAS has been tested--Raleigh-Durham Airport, Wallops Flight Center, and Manassas Airport. One often sees clutter patches so severe that the target detection module effectively blanks out these areas, thereby creating blind spots in the radar coverage. Close-in clutter is an ever-present problem and, in the absence of purposeful blanking, has the tendency to produce meandering tracks. EMI and precipitation often lead to false tracks when either is present. Finally, and one of the most pernicious problems, are ground vehicles operating on or near a field; in spite of their generally

slow speeds compared to aircraft, they nevertheless produce tracks which are difficult to distinguish from aircraft tracks.

To address these problems, the TWS module exploits additional characteristics of the targets of interest to APAS, namely general-aviation aircraft. It is noted that the development to this point depends only mildly on precise target characterization. Indeed, the only instance in which such considerations arise is in the construction of a target dynamical model for use in the tracking filter. The assumptions made on this model are the following: (1) vertical speed is generally negligible compared to horizontal speed; and (2) in-track acceleration is negligible. As a result, the model effectively excludes only high performance military aircraft and missiles, neither of which is of interest anyway.

The additional model assumptions now proposed are motivated by the observation that aircraft behavior at a busy, uncontrolled airport is not chaotic. Pilots tend to use the announced active runway; they enter and fly the traffic pattern properly; and they generally observe proper flight procedures. An aircraft observed flying toward an airport can be expected either to bypass the airport, overfly it, or enter the traffic pattern. Similarly, an aircraft observed taking off can be expected to depart the area or to enter the traffic pattern for touch-and-go practice.

The means by which this purposeful behavior is reflected in the TWS processing performed by APAS rest on the motion of a track class. The class of a track is intended to summarize in a single variable what the state of that track implies concerning its behavior with respect to accepted flight procedures. Doing so then allows inferences to be made concerning the future behavior of the target, inferences which can be used, first, to aid and manage the basic TWS activities of tracking and association and, second, to evaluate the tracks thereby created to qualify them for inclusion in traffic advisories.

Classification, the subject of the present section, begins with a separation of targets into two general types: those outside the pattern area and those inside. Generally speaking, targets outside the pattern area create fewer problems for APAS because of the obviously lower aircraft density and the generally less severe clutter problems. A target then, for which

$$d > D_p ,$$

where D_p is the radius of the pattern area (Section 2.2.2), is put into one of the three classes: arrivals, for which

$$\cos(\psi - \theta) \leq -1/\sqrt{2} ;$$

departures, for which

$$\cos(\psi - \theta) \geq 1/\sqrt{2};$$

and flybys, for which

$$-1/\sqrt{2} < \cos(\psi - \theta) < 1/\sqrt{2}.$$

The class, c , for a departure is defined numerically to be -11; that of an arrival, 11; and that of a flyby, 12.

Targets inside the pattern area, on the other hand, also fall into three main categories: above pattern altitude; in the pattern or departing; or neither. The first category, above pattern altitude, is the class of flyovers for which

$$h > H_p$$

and is assigned the value $c = -2$. The last category, the class of wanderers, is assigned $c = -1$; it is this class of targets which presents the greatest problems, since they can variously represent arriving aircraft intending to enter the pattern, aircraft flying through the pattern, ground vehicles, or close-in clutter. The pejorative name given to this class is motivated by the observation that such targets are usually of the last two types.

The final main category of targets inside the pattern area fall into twelve classes which represent an expansion of the scheme used to classify aircraft in a traffic advisory. What are labelled departures in an advisory are assigned the class $c = -10$ in the TWS module and called runway departures. Likewise, aircraft lumped together according to pattern leg in an advisory are called pattern targets and are assigned the classes $c = 1, \dots, 10$ according to the indicated table. It is seen from the following table that the determination of the class of a pattern target first uses the traffic advisory categorization for the target and then adjusts it depending on whether or not an additional constraint is met. The objective of this finer characterization is to identify targets (presumably aircraft) which may be about to turn from one pattern leg to the next.

TRAFFIC ADVISORY CATEGORY	TWS CLASS	ADDITIONAL CONDITIONS	CLASS (C)
Upwind	Upwind		10
Upwind	Turning Crosswind	$x \geq L_R/4$	9
Crosswind	Crosswind		8
Crosswind	Turning Downwind	$y \leq 0$	7
Downwind	Downwind		6
Downwind	Turning Base	$x \leq -L_R/2$	5
Base	Base		4
Base	Turning Final	$y \geq -950$	3
Final	Final		2
Final	Short Final	$x > -(L_R/2 + 800)$	1

Classification is implemented in the TWS module by including additional variables in a track data block, namely estimated class, c^* , derived from the current estimated state, X^* , and predicted class, $c^{\wedge}(n)$, derived from the current predicted state, $X^{\wedge}(t^{\wedge}(n))$. For reasons intimated above and outlined in Sections 3.1.1 and 3.1.2 each predicted class $c^{\wedge}(n)$, $n=1,2,3,\dots$, of a track is compared with its estimated class c^* . If an unallowed class transition occurs, that track is henceforth deleted from traffic advisories. Likewise, as part of updating track, the new estimated class $c^*(L)$, is compared with the last estimated class, $c^*(L-1)$. If an unallowed transition occurs in this case, the track is terminated.

3.1.3.4 Management. The function of management in the TWS module involves overseeing the initiation, evolution, and eventual termination of tracks, both as individual entities and as a group. Included also in the management function are the tasks of monitoring incoming data and coordinating the various activities of the module in processing these data. This latter group of tasks, although obviously critical to successful operation, are sufficiently mechanical that only the oversight activities are discussed here.

The first of the management tasks to be described concerns track termination. To do so requires that the notion of missing a target be defined. It is recalled from Section 3.1.3.2 that a mechanism exists for determining when a target being tracked has been bypassed, i.e., the progression of processing sectors passes over the current association sector of the track without an update of the track occurring. This circumstance is almost always caused by a failure to detect the target; the other possibility, associating the target report to the wrong track, occurs much less frequently. In general, however, it is not accurate to assert that the target has

suffered a missed detection when it is merely bypassed; the target may not be located in the active elevation beam and its detection should not be expected.

A miss, therefore, is said to occur when a track is not updated during an entire radar survey. The mathematical criterion for a miss is the following: If $X^{\wedge}(t^{\wedge}(n))$ is the current predicted state of the track and $\phi^{\wedge}(t^{\wedge}(n))$ the corresponding elevation, then one determines the elevation beam whose boresight elevation angle differs the least from $\phi^{\wedge}(t^{\wedge}(n))$. If this predicted observation beam is active and the target is bypassed, then a miss is said to have occurred. This criterion provides for the possibility that a track might pass from one beam to another and is somewhat more precise than one which uses simply a scan count.

Given, then, that a target is missed, its track enters a coasting state, the duration of which is by definition equal to the number of consecutive misses it has suffered. Once having been placed in a coasting state, the track must be updated to leave. In the meantime, the track could be coasted by setting:

$$\begin{aligned}t^{\star} &= t^{\wedge}(n), \\X^{\star} &= X^{\wedge}(t^{\wedge}(n)), \\p^{\star} &= P^{\wedge}(t^{\wedge}(n)),\end{aligned}$$

and

$$c^{\star} = c^{\wedge}(t^{\wedge}(n)),$$

the point being that the target dynamical model loses credibility for times greatly exceeding the latest update time t^{\star} . (Recall that the model is based on a series expansion.) In point of fact, however, the validity of the model is not a critical issue in the TWS module and coasting a track as indicated creates technical problems in implementation. For this reason, a missed target simply leads to incrementing the coasting counter; otherwise, the usual prediction procedure is followed.

Once a track has begun coasting, it becomes a candidate for termination. In general, one wishes to terminate a track--either one on a valid target which has left the terminal area or one caused by clutter--as promptly as possible. Since a coasting state may be produced, however, simply by one or more missed detections or by an aircraft overflying the radar and dropping out of radar coverage, some coasting should be permitted. The approach adopted in APAS is generally to allow a maximum coast count of five. If this figure is exceeded, the track is terminated. The sole exception to this rule concerns overflights: a very firm track, one for which L , the length count, is no less than 20, will be coasted as long as its predicted elevation $\phi^{\wedge}(t^{\wedge}(n))$

exceeds 26^0 , a nominal upper elevation limit on the radar coverage.

It must be emphasized that a coasting track is a potential source of trouble: the uncertainty in its predicted position grows so large that eventually almost any observation will associate with the track. For this reason it is prudent to terminate coasting tracks before their coasting counters exceed the standing threshold of five if good reason for doing so can be found. In this regard, two subsidiary termination criteria are employed in the TWS module. The first is obvious: if the predicted range of a target $r^{\wedge}(t^{\wedge}(n))$, exceeds the radar range, R_{MAX} , then the track is terminated; the target, whether an aircraft or not, is presumed to have departed that terminal area. The second criterion is more complicated and involves a comparison of the estimated class of a track, c^* , with its predicted class, $c^{\wedge}(t^{\wedge}(n))$. If in this comparison an invalid transition is observed, the track is terminated. Figure 3-5, developed in testing APAS, displays the class of transitions prohibited by the system. Note, in particular, that a track classified as short-final or runway-arrival, $c^*=0$ or $c^*=1$, is not allowed to transition to a runway departure, $c=-10$. Coasting this type of track--allowing for the possibility that it might be a touch-and-go aircraft--is so prone to uncorrect associations that it is preferable to terminate the track and, in the case of a touch-and-go, hope to reestablish it quickly rather than coast it. Close-in clutter simply creates too much room for error to do otherwise.

A second, equally broad class of termination criteria not related to coasting is based on perceived pathological behavior in a track. In this case, one seeks to identify tracks caused by clutter or jamming on the basis of the erratic behavior typically displayed by such tracks. (Note that neither ground vehicles nor stationary reflectors necessarily produce tracks falling in this particular category; indeed, it is more advantageous to maintain track on, say, a water tower than continually to terminate and reinitiate track on such a target.) Generally speaking, erratic behavior manifests itself best in the estimated velocity of a target. Thus, in the first place, any track whose wind-adjusted speed, i.e., the quantity

$$[(x^* - x_w)^2 + (y^* - y_w)^2 + z^{*2}]^{1/2},$$

exceeds 200 knots is terminated. Although this criterion may terminate valid tracks produced by en route aircraft overflying a terminal, such aircraft are typically at an altitude which would make their inclusion in a traffic advisory improper anyway.

A check on track heading similar to that above on speed is, of course, meaningless: any given track heading is certainly possible. Rather it is changes in heading which are

		Entry Class																
		-11	-10	-2	-1	0	1	2	3	4	5	6	7	8	9	10	11	12
Exit Class	-11		x			x	x											x
	-10					x	x	x	x	x	x					x	x	
	-2		x			x	x											x
	-1		x			x	x											x
	0	x	x	x	x			x	x	x	x	x	x	x	x	x	x	x
	1	x	x	x	x				x	x	x	x	x	x	x	x	x	x
	2		x							x	x	x	x	x	x	x	x	
	3		x			x					x	x	x	x	x	x	x	
	4		x			x	x					x	x	x	x	x	x	
	5		x			x	x	x					x	x	x	x	x	
	6		x			x	x	x	x					x	x	x	x	
	7					x	x	x	x	x					x	x	x	
	8					x	x	x	x	x	x					x	x	
	9					x	x	x	x	x	x	x						x
	10					x	x	x	x	x	x	x	x					x
	11		x			x	x											
	12		x			x	x											

Figure 3-5. - Invalid class transitions.

significant, erratic behavior being typically associated with inexplicable heading changes. Unfortunately, no check which involves purely heading estimates and which both performs well and is straightforward to implement has been forthcoming. In its place is used an indirect check on heading changes based on changes in pattern class. Thus any track which exhibits an invalid transition, in the sense of Figure 3-5, estimated class, $c^*(n-1)$, to another, $c^*(n)$, is terminated.

Taken as a group, the termination criteria just described--those concerned with coasting and deviant behavior--serve to reduce the incidence of false tracks in APAS to acceptable levels. In spite of their general comprehensiveness, it usually proves desirable to supplement them with additional criteria which are specific to a given airport and its unique clutter environment. Such special-purpose criteria, which soon suggest themselves in testing, are not described here.

The second facet of the track oversight function to be discussed is track splicing, by which is meant an extension of the fundamental observation-to-track association process described earlier. It often happens that an aircraft being tracked passes through a dense clutter patch, during which time it goes undetected, and then re-emerges where it is once again detected. If, during the blackout, the aircraft holds to a constant course, its track is typically coasted successfully and is properly updated when the first detection following the blackout is received. If, however, the aircraft executes a maneuver during the blackout, say a pattern turn, then its coasted track may deviate from its actual track to such an extent that resumption of successful tracking becomes impossible. Instead, the old track on the aircraft continues to be coasted and a new track on the same aircraft is initiated.

To cope with this problem, it becomes desirable to institute a higher order association process in which not observation, but rather tracks, are associated with tracks, a process termed here splicing. If, therefore, a target is predicted to be in the pattern, i.e., $0 \leq c^*(t^*(n)) \leq 10$, and if, furthermore, its track is being coasted, then the track becomes a candidate for splicing. Whenever a newly formed track enters a reasonably firm status, i.e., its length L satisfies

$$L \geq 3,$$

and this track is also classified as in the pattern,

$$0 \leq c^* \leq 10,$$

it is checked against all splicing candidates to see if it might be the continuation of one of them. To be declared so, one must have: (1) a forward progress in class,

$$c_i^*(t^*(n)) - c_j^* \leq 2,$$

where the subscript i denotes the splicing candidate and j the new track; and (2) a reasonably small jump in position, i.e.,

$$\frac{[(x_i^* - x_j^*)^2 + (y_i^* - y_j^*)^2]^{1/2}}{t_i^* - t_j^*} \leq 4 q_i^*,$$

where

$$q_i^* = [\dot{x}_i^{*2} + \dot{y}_i^{*2}]^{1/2}$$

is the estimated speed of the splicing candidate. If these conditions are satisfied, track i is terminated and the length of track j increased,

$$L_j \leftarrow L_j + L_i,$$

to enhance its firmness.

The final aspect of track management to be developed here concerns the possibility that tracks might be purposefully steered to enhance both the tracking and the observation-to-track association processes. The basis for such steering, as might be anticipated, is the track classification procedure, there being no justification otherwise for assuming anything except straight-and-level flight. In particular, for targets classified as being in the pattern, i.e.,

$$0 \leq c^* \leq 10$$

or

$$c^* = -10 \text{ (runway departure),}$$

there is reason to believe either that the target (aircraft) is executing some maneuver typical of the pattern class or that it may be about to do so. One is thus led to adjusting the two controls present in the target dynamical model--turn rate, ω , and rate of climb, ρ --to effect the proper target maneuver.

Steering is implemented in APAS in one of several distinct forms. The first can be thought of as a rigorous observation of the tracking filter methodology; that is, in predicting target motion for purposes of both observation-to-track association and track

updating, one sets both turn rate, $\omega^{\wedge}(t^{\wedge}(n))$, and rate of climb, $\rho^{\wedge}(t^{\wedge}(n))$, in calculating $X^{\wedge}(t^{\wedge}(n))$ and maintains them at this value throughout. This subsumes the usual (non-steering) situation, with both steering controls set at zero. A target, however, observed to be turning final or on final,

$$1 \leq c^* \leq 3,$$

is presumed to be descending. If turning final, $c^* = 3$, one sets its predicted rate of climb

$$\rho^{\wedge}(t^{\wedge}(n)) = -16 \text{ ft./sec.}$$

for all subsequent predictions; in the other two cases, one sets

$$\rho^{\wedge}(t^{\wedge}(n)) = -25 \text{ ft./sec.}$$

Although admittedly a rough prediction, these rates of climb (descent) are adequately reflected actual aircraft motion. On the other hand, if a target has class $c^* = -10$ and is presumably climbing out, one sets

$$\rho^{\wedge}(t^{\wedge}(n)) = 16 \text{ ft./sec.}$$

to create an ascending track.

A second, somewhat more severe form of steering, involves simply overriding the output of tracking filter. Thus, if a target is classified as upwind, $c^* = 10$, one sets ψ^* equal to the runway heading and then adjusts \dot{x}^* and \dot{y}^* to realize this new heading, holding horizontal speed q^* fixed. An analogous procedure, with the appropriate heading, is followed for aircraft on crosswind ($c^* = 8$), on downwind ($c^* = 6$), on base ($c^* = 4$), and on final ($c^* = 2, 1$, or 0). The effect of these adjustments is to force targets to follow a more orderly trajectory in the pattern.

The third form of target steering, the most drastic, consists of actually flying a target observed to be in the pattern. If, for example, a pattern target classified as turning enters a coasting state, it is possible actually to steer the target around a pattern turn. If done properly, i.e., if a turn has been predicted accurately, it then typically proves unnecessary to depend upon splicing to match up the otherwise coasted track with a newly formed track. Mathematically, this form of target steering is implemented by choosing the steering controls to produce the desired target motion and then to set the estimated state to the predicted state:

$$\begin{aligned}t^* &= t^{(n)} \\x^* &= x^{(t^*(n))} \\c^* &= c^{(t^*(n))}\end{aligned}$$

and

$$p^* = p^{(t^*(n))},$$

where n refers to whatever prediction stage is chosen to reset the track.

The only time this steering is performed in APAS involves aircraft on base ($c^* = 4$ or 3). Such aircraft, in the process of losing altitude, oftentimes go undetected for some period of time. Since aircraft turning final are of some concern to departing aircraft waiting to take the active runway, it is important to announce their presence. For this reason, APAS assumes that targets with firm tracks shown to be on base do, in fact, turn final if they enter a coasting status.

3.1.3.5 Evaluation. The final function of the TWS module, the evaluation of tracks, is very similar in form to the termination of tracks discussed in the preceding section. In both cases one is passing judgement on a track; however, whereas in termination an unfavorable judgement leads to the complete elimination of a track, in evaluation it simply eliminates mention of the track (i.e., the target it represents) in a traffic advisory. Together, termination and evaluation act as a two-stage sieve on tracks: the first attempts to identify and filter all spurious tracks caused by such things as vegetation, precipitation, and radio interference; the second attempts to select from all remaining tracks those which truly represent aircraft.

One can argue, of course, that track evaluation is logically a function of the voice response unit which, after all, is responsible for formatting traffic advisories. Although valid in principle, this statement ignores the intimate connection of evaluation with TWS processing; it proves far more convenient from an implementation standpoint to evaluate tracks in the TWS module.

In general, the evaluation of a track is based on its length counter, L , which in turn reflects the firmness of the track. Thus a track whose length does not exceed three is treated as tentative and is judged not yet suitable for reporting. In this regard, further discussion of the length counter is in order. Although defined originally (for simplicity) as the number of times a track has been updated, including its initiation, in point of fact this counter is manipulated in the TWS module more in line with its new interpretation as firmness indicator. In this way, one obtains a more general-purpose counter which finds use both in observation-to-track association and evaluation.

Initialized at one when a track is initiated, the length counter is incremented by one each time a track is updated until it reaches a maximum value of twenty, at which point it is held. Adjustments, however, may be made in the counter any time certain conditions on the speed of the target and its class are met. Thus a target classified as a runway departure ($c^* = -10$) has its length counter immediately set to seven to insure that the target is reported. Similarly, although not so pronounced, is the immediate increment firm track (1) whose new estimated horizontal speed is observed to differ by less than ten knots from its average estimated horizontal speed or (2) whose strength exceeds 450. Here average estimation speed, \bar{s}^* , is set to s^* at the second track update and computed by

$$\bar{s}^* \leftarrow 0.8 \bar{s}^* + 0.2 s^*$$

after each subsequent update.

On the other hand, a firm track whose average estimated speed \bar{s}^* is less than 40 knots has its length counter held or reset to three, a non-reporting status, to reflect what is presumably a stationary target or perhaps a ground vehicle. Also afforded this treatment is a firm track whose newly estimated horizontal speed differs by more than 60 knots from its average estimated speed. And, finally, overflights, tracks for which $z^* > 3,000$ feet, also have their length counter held at three.

Further conditions concerning the manipulation of the length counter could be stated, but as in the case of track termination criteria, to do so here would detract from the general overview of APAS provided in the present report. The point should be clear that the length counter provides a convenient vehicle to control fully track evaluation and, hence, reporting.

3.2 Weather Data Unit

Serving as the second main sensor unit of APAS, the weather data unit is responsible for measuring, processing, and displaying weather data. As part of these activities, it also selects and displays the active runway, using for this purpose processed wind data as well as operator inputs through the operator control panel. In addition to providing data to the voice response unit for inclusion in airport advisories, the weather data unit relays weather data (i.e., wind conditions) to the tracking data unit for use in tracking aircraft.




Figure 3-6. - Weather data unit.

3.2.1 Overview

The weather data unit consists of the three modules shown in Figure 3-6: a data acquisition module which acquires and stores raw weather data samples, a data processing module which processes and displays weather data, and a runway selection module which selects and displays the active runway. The first module operates on a more or less continuous basis, maintaining in the microcomputer memory a record of raw data samples measured by the five weather sensors (wind direction, wind speed, barometric pressure, temperature, and dewpoint). The latter two modules are tightly coupled in their respective activities which, in turn, are equally tightly coupled to the broadcasting of airport advisories, the intent being (1) to process weather data on a regular basis and thereby to broadcast up-to-date weather information; and (2) to coordinate closely the task of actually selecting an active runway with that of announcing the active runway.

To achieve the necessary coupling, the timer used to schedule the broadcasts of airport advisories is adapted for use as a timer to control the processing of weather data and the selection of an active runway. Thus the activities of the data processing and runway selection modules are controlled directly by the voice response unit, relieving the weather data unit of any independent control responsibilities and, at the same time, ensuring timely data in airport advisories.

Suppose, then, that the system operator requests T seconds to elapse between successive airport advisories. Define a threshold count -

$$N = 1 \quad \text{if} \quad T \geq 60 \text{ sec} ;$$
$$= 60/T \text{ otherwise .}$$

It is implicitly assumed in the latter case that T is 15, 20, or 30 seconds, corresponding respectively to 4, 3, or 2 airport advisories per minute. With respect to this new parameter, weather data processing and runway selection are controlled as follows. At system startup, the timer controlling the broadcast of airport advisories is initialized at zero (see Section 3.3.5), as in an internal counter which is to control the activities in question. When the advisory timer reaches T , signifying an airport advisory broadcast is called for, the internal counter is incremented and compared to the threshold count N . If the counter has not reached N , then the system immediately proceeds to the indicated advisory broadcast. If, however, the counter is equal to N , then APAS first activates the data processing module to process weather data and then the runway selection module to select an active runway, at the conclusion of which the counter is reset to zero. The system then commences broadcasting with an airport advisory which reflects the newly

processed weather data and active runway selection.

The indicated approach realizes the two desired goals: (1) weather data is processed at least once every two minutes, (it is assumed that T does not exceed 120 seconds) but no more often than once per minute; and (2) the delay between announcing that a given runway is to become active and the announcement that it is in fact active is guaranteed not to be less than one minute.

As described, the operation of the data processing and runway selection modules of the weather data unit is seen to be totally nonautomatous. They thus appear simply as subroutines within the microcomputer, to be called at will by the voice response unit. By contrast, the data acquisition module operates in just the opposite fashion as an interrupt handler, requesting and accepting a continuous stream of data from the weather sensors.

3.2.2 Data Acquisition Module

As just stated, the data acquisition module operates continuously and autonomously, basic timing for its operation being provided by the analog-to-digital conversion element. Every second, this element interrogates one of the sensors, the interrogations proceeding sequentially through the entire collection of sensors and then repeating. The stream of sensor readings so obtained are converted to digital form, transformed into floating-point numbers with the proper dimensions, and stored for later processing. If a sensor has been placed in a nonoperational status (see below), then its readings, after dimensioning, are displayed on the operator control panel.

3.2.2.1 Weather sensors. The data acquisition module is configured to accommodate up to eight weather sensors. cursory specifications on the five sensors currently implemented in the system are given in Figure 3-6.

The analog weather data samples are quantized using twelve bits of precision, a relative measurement error of one part in 4096. Combining this factor with the measurement ranges of Table IV provides the quantization errors shown in the table. In general, these errors are seen to be dominated by the intrinsic sensor measurement errors.

3.2.2.2 Sampling. Given that one of the weather sensors is interrogated every second and that there are five sensors, it would follow that a sample from a given sensor is acquired every five seconds. That this is not precisely the case results from the special treatment afforded the pair of wind sensors: in order to obtain a true reading of wind velocity, a vector quantity, these sensors are sampled virtually simultaneously. It follows, then, that the basic sampling period of a given sensor is four seconds, a period sufficiently short to provide adequate measurements on the meteorological variables of interest.

TABLE IV. - METEOROLOGICAL SENSOR SPECIFICATIONS

SENSOR	RANGE	SENSOR ERROR	QUANTIZATION ERROR
Wind Speed	0 to 87 kt	± 0.2 kt	± 0.01 kt
Wind Direction	0 to 540 deg	± 1 %	± 0.07 deg
Barometric Pressure	27.8 to 32.2 in of Hg	± 0.006 in of Hg	± 0.0005 in of Hg
Temperature	-22 to 120° F.	$\pm 0.3^\circ$ F.	$\pm 0.02^\circ$ F.
Dew Point	-22 to 120° F.	$\pm 2.7^\circ$ F. ($\pm 5.4^\circ$ F. below -10° F.)	$\pm 0.02^\circ$ F.

Weather data samples are written into computer memory using a circular list technique with sixteen storage locations allocated to each sensor. Implied, therefore, is the presence in memory at all times of the last sixteen samples taken from each sensor (following, of course, a startup transient of roughly four minutes during which the individual storage locations are filled for the first time). Having available readings from each sensor which span approximately one minute allows the various fault-detection and averaging procedures discussed below.

3.2.3 Data Processing Module

The basic function of the data processing module is to convert the series of data records from each of the sensors--spanning the past minute--into summary statistics suitable for broadcast and display. An important preliminary function of the module, however, is setting the operational status of the weather sensors. This it does in two ways: examination of the operator control panel for a switch setting indicating that the assigned sensor has been declared nonoperational by the system operator, and examination of the individual data records for pathological behavior indicative of a faulty sensor. The precise processing and fault detection afforded each of the sensors is outlined below, differing as they do from sensor to sensor.

So long as a sensor remains operational, the same sequence of steps is repeated each time data from the sensor is processed:

- (1) check for an operator-declared fault, and, if so, terminate further processing;
- (2) check for pathological data and, if so, declare a fault and terminate further processing;
- (3) process and display the data.

In all cases, the output of the data processing module is made available to the voice response unit for inclusion in airport advisories.

Once a sensor becomes nonoperational, for whatever the reason, the data processing module monitors the operator control panel for a change in the status switch of that sensor. If it notes a change in which the switch, previously in the faulted position, now appears in the opposite position, then the module returns the sensor to operational status. This approach allows the system operator to clear an APAS declared fault through his declaring a manual fault and then clearing that fault, a procedure which reflects a repair--real or simulated--to the faulty sensor.

3.2.3.1 Temperature and dewpoint. The two meteorological variables of temperature and dewpoint are treated identically. Automatic fault detection simply involves a check on the last recorded value to ascertain whether it lies in a prespecified range set by

the temperature extremes characteristic of the geographic region in which APAS happens to be located. If the last reading lies outside this range, it is assumed that a sensor malfunction (typically an open- or short-circuit) has occurred and the sensor--temperature or dewpoint as the case may be--is declared non-operational. If, on the other hand, the reading passes the fault test, it is designated the processed value of the variable in question and is made available to the voice response unit for announcement. Simultaneously, this value is displayed on the operator control panel, rounded to the nearest degree Fahrenheit and with the appropriate sign.

3.2.3.2 Altimeter. Somewhat more complicated is the processing afforded the barometric pressure sensor. As in the case of temperature and dewpoint, the last recorded pressure reading must also fall in a prespecified range (nominally 29.00 to 31.00 in of Hg) to prevent a faulty sensor diagnosis. In addition, it is required that the deviation of this latest reading from that taken one minute earlier not exceed 0.05 in of Hg, the intent in this test being to detect sporadic behavior indicative of a faulty sensor but uncharacteristic of barometric pressure.

Given a latest pressure reading passing these tests, it is converted into an altimeter reading according to the following formula:

$$p_a = p \left[1 + 6.87918 \times 10^{-6} H_b \left(\frac{p_0}{p} \right)^n \right]^{1/n},$$

where

p_a = altimeter reading (in of Hg)

p = measured barometric reading (in of Hg)

H_b = elevation of pressure sensor (ft)

p_0 = 29.921 in of Hg

n = 0.190284 .

The parameter H_b specifying airport elevation is entered into the system at startup. Given the numerical difficulties in evaluating the above formula, the following approximations are used instead:

$$\hat{p} \triangleq \left(\frac{p_0}{p} \right) ,$$

then

$$\hat{p}^n \approx \alpha_0 + \alpha_1 \hat{p} + \alpha_2 \hat{p}^2$$

where

$$\alpha_0 = 0.73235$$

$$\alpha_1 = 0.34485$$

$$\alpha_2 = 0.07720 .$$

Moreover, if

$$\hat{H}_b = 6.87918 \times 10^{-6} H_b \hat{p}^n$$

and

$$f = (1 + \hat{H}_b)^{1/n}$$

then

$$f \approx \beta_0 + \beta_1 \hat{H}_b + \beta_2 \hat{H}_b^2 + \beta_3 \hat{H}_b^3 ,$$

where

$$\beta_0 = -0.999994$$

$$\beta_1 = 5.25525$$

$$\beta_2 = 11.1973$$

$$\beta_3 = 12.1370 .$$

Finally

$$p_a = p \hat{f} .$$

The altimeter reading is displayed on the operator control panel to the standard 0.01 in of Hg. It should be noted that the processing just described for converting barometric pressure to altimeter is performed only in the no-fault situation. If, therefore, the pressure sensor is declared nonoperational, the operator control panel display reverts to a raw barometric pressure reading.

3.2.3.3 Wind direction and speed. Easily the most elaborate processing performed by the data processing module is accorded the wind sensor data, both in fault detection and in the computation of the wind statistics of interest: average speed, average direction, and gusts (or maximum speed). In the case of fault detection, one must ascertain deviant behavior in a series of measurements that as a rule display rapid and extreme variations which are totally atypical of the sensors discussed above. This same variability requires averaging over a series of measurements to obtain a useful wind characterization, an operation which is largely superfluous with barometric pressure and temperature.

The fault detection approach implemented in the DPM is based on the fundamental assumption that static wind data is indicative of a sensor fault, i.e., that a series of consistently high or consistently low readings from a sensor signify an open - or short-circuit. (In contrast, a single abnormally high or abnormally low reading from one of the previous sensors may be taken as a fault indication). Blind application of this assumption, however, will fail to cope successfully with the calm wind situation. The fault detection scheme used, therefore, attempts to recognize any static behavior which is not attributable to calm winds.

To this end, the wind sensor readings obtained over the last minute (fifteen measurements per sensor) are scanned to produce three statistics:

- (1) a count of the number of times wind direction measurements are less than 1° or greater than 439° *
- (2) a count of the number of times wind speed measurements exceed 85 knots
- (3) the maximum wind speed measurement.

If the last statistic, maximum wind speed, proves to be less than one knot, a calm wind situation is declared and fault detection processing is terminated. Otherwise the two

*Recall that the direction sensor is the 0° to 540° type.

counts referred to in (1) and (2) above are compared to fifteen (their maximum value), and if either equals this maximum, the corresponding sensor is declared faulty.

The final step of fault processing reflects the critical nature of wind direction in APAS (and in aviation in general). In the case of a manually-declared fault, the two wind sensors are tracked as a pair and both (or neither) becomes nonoperational; this approach is deemed appropriate in what is ostensibly a sensor repair situation. In the case of automatic fault detection, however, the wind direction sensor is maintained in an operational status, if such is indicated, irrespective of the status of the wind speed sensor. The system, therefore, will continue to display and broadcast average wind direction whether or not an average wind speed is available. The converse is not true however. A fault declared automatically in the wind direction sensor is treated the same as a fault declared manually: both sensors become nonoperational.

It is noted that a fault in the wind speed sensor alone precludes determination of a calm wind situation. In this special situation, fault checking on the wind direction sensor is allowed to continue under the more demanding assumption of a non-calm situation. It is possible, therefore, but unlikely, that calm winds could cause an erroneous fault declaration in the wind direction sensor.

The processing accorded the wind data recognizes the vector nature of wind velocity. To compute average wind velocity, therefore, the fifteen wind velocity measurements used in fault processing are first converted to rectangular coordinates, averages of the respective components computed, and the resulting average wind velocity converted back to polar coordinates. In the case of a wind speed sensor fault, computation of an average wind direction proceeds in the same way under the assumption of a constant wind speed of one knot.

The indicated conversion from polar to rectangular coordinates requires computation of the sine and cosine of the wind direction. The computation is performed by a general-purpose routine which computes:

$$x = r \cos(\theta)$$

$$y = r \sin(\theta)$$

given r and θ . The routine proceeds in the obvious way using trigonometric identities and the following approximations, valid for $0 \leq \theta \leq \pi/2$:

$$\cos(\theta) \approx \alpha_1 \theta^2 + \alpha_2 \theta^4,$$

where

$$\alpha_0 = 1.0$$

$$\alpha_1 = -0.4963$$

$$\alpha_2 = 0.03705 ,$$

and

$$\frac{\sin(\theta)}{\theta} \approx \beta_0 + \beta_1 \theta^2 + \beta_2 \theta^4$$

where

$$\beta_0 = 1.0$$

$$\beta_1 = -0.16605$$

$$\beta_2 = 0.00761 .$$

The inverse transformation from rectangular to polar coordinates is performed likewise, computing:

$$r = (x^2 + y^2)^{1/2}$$

$$\theta = \tan^{-1} \frac{y}{x} ,$$

with θ placed in the appropriate quadrant, according to the following approximations:

$$(1 + \alpha^2)^{1/2} \approx \gamma_0 + \gamma_1 \alpha^2 + \gamma_2 \alpha^4 + \gamma_3 \alpha^6, \quad |\alpha| \leq 1 ,$$

where

$$\gamma_0 = 1.0$$

$$\gamma_1 = 0.495703$$

$$\gamma_2 = -0.102980$$

$$\gamma_3 = 0.021515;$$

and

$$\tan^{-1}(\rho) \approx (\delta_0 + \delta_1 \rho^2 + \delta_2 \rho^4) \rho, \quad |\rho| \leq 1,$$

where

$$\delta_0 = 0.994949$$

$$\delta_1 = -0.287061$$

$$\delta_2 = 0.078037$$

Computation of wind gusts requires a somewhat different approach, since it is defined as the maximum wind speed over the preceding five minutes, an awkward time interval compared to the one-minute storage of wind data and the time interval, from one to two minutes, separating successive periods of weather data processing. (See the preceding Section 3.2.1). To avoid the additional storage and/or control required to adhere strictly to its definition in the computation of wind gusts a simple approach is adopted which selects--properly--a five-minute maximum but which uses wind speed data over a time interval which may vary from five to ten minutes, depending on the rate at which the data processing module is invoked. Thus, every time it is called, the module selects the maximum in push-down stack five slots deep, and takes as the wind gust value the maximum of the five wind speeds in the stack. (It should be noted that the wind speed stack is initialized to zero at system startup; in a fault condition, a zero value is pushed on the stack, forcing the stack to all zeros if the fault persists long enough).

In an operational status, wind direction is always displayed on the operator control panel rounded to the nearest 10° , $0^\circ - 360^\circ$; wind speed and gusts are displayed to the nearest knot. In a fault condition, however, wind direction appears to the nearest degree, $0^\circ - 540^\circ$, and wind speed and gusts are made identical.

3.2.4 Runway Selection Module

The preceding Section 3.2.1 indicates the approach taken in APAS in coordinating the selection of an active runway with the processing of weather data and the announcement of runway status in airport advisories. The runway selection module is invoked at a rate not exceeding once per minute to assess the latest wind conditions and current runway

status and, if necessary, to choose an alternative runway to serve as the active. Should a new active runway be selected, then in the forthcoming airport advisory, and in all succeeding advisories until the runway selection module is next invoked, this runway will be announced as the runway about to become active ("CHANGING TO..."). In this general context, then, the runway selection module plays strictly a subservient role: when invoked, it simply selects a preferred candidate for the active runway and duly reports its choice. All timing and control responsibility concerning runway changes resides in the voice response unit, thereby guaranteeing coordinated runway status announcements.

The operation of the runway selection module is predicated in the maintenance at all times of a pair of runway identification numbers, the first referring to the so-called current preferred runway (CPR) and the second to the next preferred runway (NPR). As part of the system startup procedure, the NPR is set to zero (0), the identification number of an imaginary runway (by convention). Following startup, the module, whenever it is invoked, routinely promotes the NPR to CPR status and proceeds to select a new NPR which is to designate that runway representing the best choice to serve as the active runway under the currently prevailing conditions.

The interpretation attached at any given time to any given values of the NPR and the CPR are as follows:

- if $CPR = 0$, no active runway is designated (irrespective of NPR)
- if $CPR \neq 0$, but $NPR = 0$, an active runway is designated (CPR) but (by convention) no runway change is in progress
- if $CPR = NPR \neq 0$, an active runway is designated (CPR) and no runway change is in progress
- if $CPR \neq 0$ and $NPR \neq 0$ and $CPR \neq NPR$, an active runway is designated (CPR) but a runway change is in progress (NPR)

On the basis of this pair of runway identification numbers, then, the system as a whole may infer at any time the runway status and, in particular, broadcast and display the proper information to this effect.

Given the procedural approach just described, the main function of the runway selection module becomes the selection of the NPR. This it does in a distinct series of steps: (1) identify the available runway candidates as indicated on the operator control panel; (2) evaluate the suitability of all candidates, in the process

associating to each a single number representing the desirability of that candidate; and (3) choosing as the NPR the highest ranking candidate. Needless to say, if only one candidate runway is available (the manual mode of runway selection), this procedure, although not overly efficient, nevertheless will always rank the single candidate highest and choose it as the NPR.

The key to the whole procedure, of course, is the technique whereby the desirability of a given runway candidate vis-a-vis all other candidates is reflected in a single number. To this end, the following reward function is employed: the reward of runway candidate i is given by:

$$r_i = w_i + p_i ,$$

if i is not the CPR, and is:

$$r_i = w_i + p_i + p$$

otherwise. Here w_i is the component of the wind velocity along the runway, p_i is a premium attached to runway i which reflects its relative desirability, and p is premium attached to the current active runway. By making the premium of runway i high compared to that of runway j , the former will clearly be preferred to the latter, all else being equal. Furthermore, by increasing the active runway premium, p , one can forestall runway changes until a candidate clearly superior to the present active runway emerges, thereby preventing overly frequent runway changes.

A numerical example serves to motivate the runway selection algorithm. Consider a four runway geometry in which the headings of the respective runways are 30° , 90° , 210° , and 270° (denoted ψ_i , $i = 1, \dots, 4$, in what follows). Assume initially that none of these runways is to be preferred over any other, so that the runway premiums may be set to zero:

$$p_i = 0 \text{ kt}, \quad i = 1, \dots, 4 .$$

In the case in which no active runway has been designated, the preferred candidate to serve in this capacity is simply the one with the maximum wind component, since $r_i = w_i$ for all candidates. If the wind velocity w is expressed in terms of rectangular components u and v , then it is found that the set of wind conditions for which runway i is preferable to runway j is described by the linear inequality:


$$u[\cos(\psi_i) - \cos(\psi_j)] + v[\sin(\psi_i) - \sin(\psi_j)] \geq 0 .$$

Consideration of these inequalities leads to Figure 3-7a in which the various wind conditions favorable to the respective runways are shown.


Now suppose that runway 2 (i.e., 09) is designated the active and that the active runway premium has been chosen as 5 kt. In order for an alternative runway to become the active, it is necessary not only that the wind be favorable for this runway in the sense of Figure 3-7a but also that it be distinctly favorable by at least 5 kt. This situation is depicted in Figure 3-7b, in which it is seen that region of wind conditions for which runway 2 will remain active is enlarged compared to the previous figure and that all other regions are made correspondingly smaller.

A similar analysis is shown in Figure 3-7c and 3-7d for same runway configuration in which all runways are not judged equally desirable. In this case, it is assumed that 03 and 21 are less desirable than 09 and 27, leading one to set $p_1 = p_3 = 0$ kt. and $p_2 = p_4 = 5$ kt. Comparison of Figure 3-7c with Figure 3-7a shows, as one would anticipate, that 09/27 will always be preferred to 03/21 in light and variable conditions. It should be noted, in concluding this example, that both the individual runway premiums and the active runway premium are design parameters and must be adjusted on the basis of experience of an individual airport.


For the most part, the algorithm used by the runway selection module performs well. In spite of the active runway premium concept, however, excessively frequent changes sometimes occur in highly variable wind conditions, particularly in the case in which there are only two runway candidates. It would seem desirable, therefore, to always maintain a static runway status for some length of time greater than the one to two minutes currently realized by the system. In contemplating such a modification to the existing selection algorithm, one must, of course, balance the ability of the system to cope with variable wind conditions with its responsiveness to a true wind shift calling for a runway change.




(a) No current active



(b) Runway 03 active



(c) No current active



(d) Runway 03 active

Figure 3-7. - Runway selection algorithm.

3.3 Voice Response Unit

The voice response unit (VRU) interfaces to both the tracking data unit and the weather data unit and has as its main role in APAS the verbalizing of data provided by these two sensor units. An additional capability provided by the VRU is a means for an FBO to broadcast information of an unconstrained and transient nature (a so-called discretionary message) which serves as a supplement to airport advisory messages.

3.3.1 Overview

The form taken by the voice response unit in APAS is shown in Figure 3-8. The unit is seen to consist of four major modules: FRU timing module, message formatting module, message entry module, and speech processing module. The last of these, the speech processing module, operates in a bilateral mode. As an output device, it converts stored digital speech into an analog speech signal directed to either the speaker in the operator control panel or the VHF transmitter. The content of the output speech is in all cases controlled by the message formatting module, which translates numerical information provided by the tracking data and weather data units into word sequences according to a prescribed format. As an input device, the speech processing module accepts aural input from the system operator through an operator control panel microphone, transforms the analog speech signal into a digital representation, and stores the digitized speech for later playback as a discretionary message. This mode of operation is under direct control of the message entry module. Responsible for overall coordination of the VRU operation is the VRU timing module, which schedules the delivery of all system broadcasts and which interrupts the broadcast schedule for the entry of discretionary messages.

Although possessing considerable flexibility, the voice response unit as currently implemented produces only three types of output messages: airport advisory, traffic advisory, and discretionary message playback. The last is not to be confused with discretionary messages appended to airport advisories. Rather, they represent the final operation in the entry of discretionary messages, in which a just-recorded message is played back locally through the operator control panel speaker to allow the system operator to check it for clarity and correctness before allowing its broadcast. With reference to Figure 3-8, the playback operation is initiated by the VRU timing module after it receives confirmation from the message entry module that a new discretionary message has been recorded.

The performance requirements placed on any speech synthesis system for use in APAS are, in the first place, the obvious ones: it must create messages of the diversity

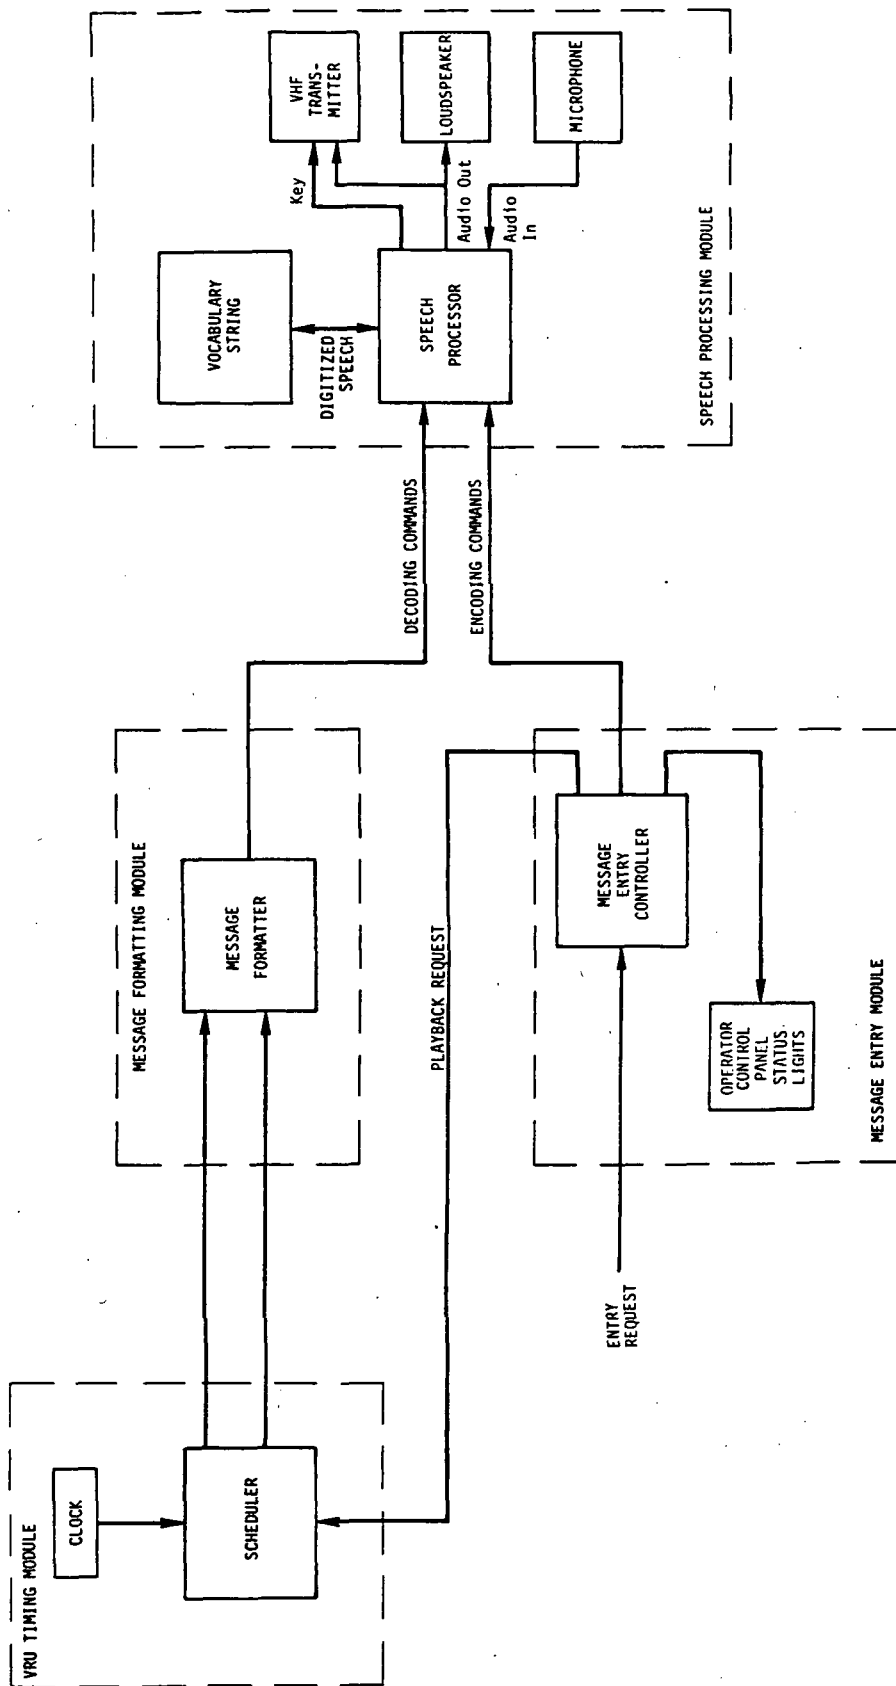


Figure 3-8. - Voice response unit (VRU).

illustrated in Section 2.2, and it must deliver these messages with synthetic speech displaying high intelligibility and good, if not superb, quality. A somewhat more subtle requirement is that the synthetic speech should possess a distinct machine accent, the intent being not to confuse pilots into believing that APAS is a person with whom they can communicate.

Based on these requirements, the voice response unit utilizes a waveform encoding technique, specifically, adaptive differential pulse code modulation (ADPCM), to create digitized versions of individual spoken phrases, the phrases being chosen so that any particular version of any of the three types of system messages can be constructed by concatenating a sequence of such phrases. In this context, a phrase may consist of a single word (e.g., "ONE"), several words (e.g., "LIGHT AND VARIABLE"), or--the limiting case--the complete sentence (s) making up a discretionary message (e.g., "CAUTION. MOVING OPERATIONS IN PROGRESS."). This particular approach provides the wherewithal to construct the variety of sentences required by airport and traffic advisories and to verbalize these sentences with high intelligibility and quality.

Although there are speech encoding techniques more economical than ADPCM in terms of required memory, the extra memory is compensated for by the simplicity of an ADPCM-decoder and the overall quality of the speech it produces. Furthermore, an ADPCM-decoder lends itself to a bilateral implementation which allows its use as an encoder of discretionary messages as well. As a subsidiary benefit, the concatenation of individually encoded phrases naturally produces a machine accent which is readily discernable but not overly objectionable.

With the exception of the discretionary message phrase, the phrases referred to above are entered into the voice response unit memory at the time of APAS startup and are not changed thereafter. They, together with any discretionary message phrase subsequently entered into the system by the FBO, form the VRU vocabulary. For this reason, and to avoid terminology problems later, these phrases are henceforth referred to as words.

The vocabulary of APAS is listed in Figure 3-9. It is seen from this figure that the choice of words comprising the vocabulary are not unique. For example, one could substitute the set of words: "AIRPORT", "TRAFFIC", AND ADVISORY." for the set of words: "AIRPORT ADVISORY." and "TRAFFIC ADVISORY." with a net savings in memory. The alternative, however, leads to a loss of speech quality in the important introduction to all advisory broadcasts and is therefore rejected.

The periods following some of the words in Figure 3-9 are indicative of a secondary feature of the VRU vocabulary, namely inflection. Consideration of the format of airport and traffic advisories reveals that certain words, if they are used at all, appear always at the end of sentences, e.g., "NOT AVAILABLE.". In the preparation of the VRU vocabulary, therefore, these words are recorded with the falling inflection characteristic of sentence

completion. The result is higher quality APAS speech and, it is believed, improved intelligibility. Some words, notably "AIRCRAFT", can appear both within and at the end of sentences; such words are recorded with neutral inflection.

3.3.2 Speech Processing Module

The earlier Figure 3-8 identifies the two generic elements of the speech processing module: a processor/memory and a collection of input/output devices, namely a VHF transmitter, a loudspeaker (and associated amplifier), and a microphone. Operation of this module proceeds on a per word basis (where word is used in the sense defined above); the message formatting module commands it to play back a given word stored in memory; or the message entry module commands it to record a given word (i.e., store the word in memory).^{*} The commands to the speech processing module always refer to individual words by their identification numbers, assigned more or less arbitrarily from the integers 1 through 250 (see Figure 3-9). By convention, the discretionary message is assigned 251 as its identification.

Appended to the vocabulary are four special words, with identification numbers 252 through 255, which are pauses (i.e., silence) of varying durations. Word 252 is a pause of duration 500 msec which begins every advisory broadcast to allow time for the transmitter key to take effect. Words 253 and 254 are pauses designed to optimize the intelligibility and quality of spoken output. The former, of duration 39.1 msec, is inserted after every word so that words are separated by periods of silence; the latter, of duration 76.2 msec, is inserted after every sentence. Both values were chosen on the basis of informal listening tests.^{*} The final word 255 is an end-of-message pause of duration 1000 msec which simply guarantees that two consecutive advisories are well separated.

Figure 3-10 is an expanded block diagram of the speech processing module. The central elements of the module are the coder which, through the appropriate switch settings, functions either as a decoder of digital speech or as an encoder of analog speech; the vocabulary store, the memory in which all digital speech resides; and the controller, which controls the coder and effects all transfers of data between coder and vocabulary store. The remaining elements, the data handler and the command file, create the interface between the speech processing module on the one hand and the message formatting and message entry modules on the other. Specifically, the data handler converts a list of words to be verbalized (or recorded), as specified by the last named modules, into a corresponding sequence of commands to the controller and posts the commands in the command file. From this buffer file the controller retrieves and executes commands in sequence.

^{*}The only word in question is, of course, the discretionary message.

<u>IDENTIFICATION</u>	<u>WORD</u>	<u>LENGTH</u>
1	AIRPORT ADVISORY.	787
2	HYDE FIELD.	565
3	SALISBURY.	475
4	MONTGOMERY COUNTY.	647
5	MANASSAS.	439
6	WALLOPS FLIGHT CENTER.	712
7	GMT	662
8	NOT AVAILABLE.	511
9	TEMPERATURE	407
10	DEWPOINT	362
11	CEILING	287
12	WIND	288
13	AT	215
14	GUSTING TO	490
15	LIGHT AND VARIABLE	686
16	ALTIMETER	383
17	ACTIVE RUNWAY	600
18	FAVORED RUNWAY	653
19	CHANGING TO	529
20	RIGHT HAND PATTERN	743
21	MINUS	326
22	ZERO	333
23	ONE	239
24	TWO	205
25	THREE	232
26	FOUR	212
27	FIVE	179
28	SIX	190
29	SEVEN	246
30	EIGHT	193
31	NINE	341
32	TEN	260
33	ELEVEN	301
34	TWELVE	320

Figure 3-9. - VRU dictionary.

<u>IDENTIFICATION</u>	<u>WORD</u>	<u>LENGTH</u>
35	CONFLICT ALERT.	627
36	HAS TRAFFIC AT	571
37	O'CLOCK.	294
38	TRAFFIC ADVISORY.	697
39	AIRCRAFT	453
40	AWAITING DEPARTURE.	650
41	ON FINAL	489
42	ON BASE.	495
43	ON DOWNWIND.	654
44	ON CROSSWIND.	633
45	ON UPWIND.	556
46	TURNING	290
47	ABOVE PATTERN ALTITUDE.	894
48	ARRIVING	370
49	DEPARTING	374
50	MILE	316
51	MILES	396
52	POINT	234
53	ONE HALF	377
54	HEADING	273
55	NORTH	289
56	NORTHEAST	407
57	EAST	245
58	SOUTHEAST	423
59	SOUTH	268
60	SOUTHWEST	456
61	WEST	280
62	NORTHWEST	567
63	ALFA	272
64	BRAVO	387
65	CHARLIE	269
66	DELTA	358
67	ECHO	289
68	FOXTROT	383
69	GOLF	268
70	HOTEL	382

Figure 3-9. - Continued

<u>IDENTIFICATION</u>	<u>WORD</u>	<u>LENGTH</u>
71	INDIA	337
72	JULIETT	327
73	KILO	343
74	LIMA	305
75	MIKE	204
76	NOVEMBER	376
77	OSCAR	351
78	PAPA	274
79	QUEBEC	284
80	ROMEO	358
81	SIERRA	351
82	TANGO	328
83	UNIFORM	468
84	VICTOR	372
85	WHISKEY	295
86	XRAY	362
87	YANKEE	323
88	ZULU	306
251	DISCRETIONARY MESSAGE	4201
252	BEGINNING-OF-MESSAGE PAUSE	595
253	END-OF-WORD PAUSE	32
254	END-OF-SENTENCE PAUSE	64
255	END-OF-MESSAGE PAUSE	1680

Figure 3.9 - Concluded


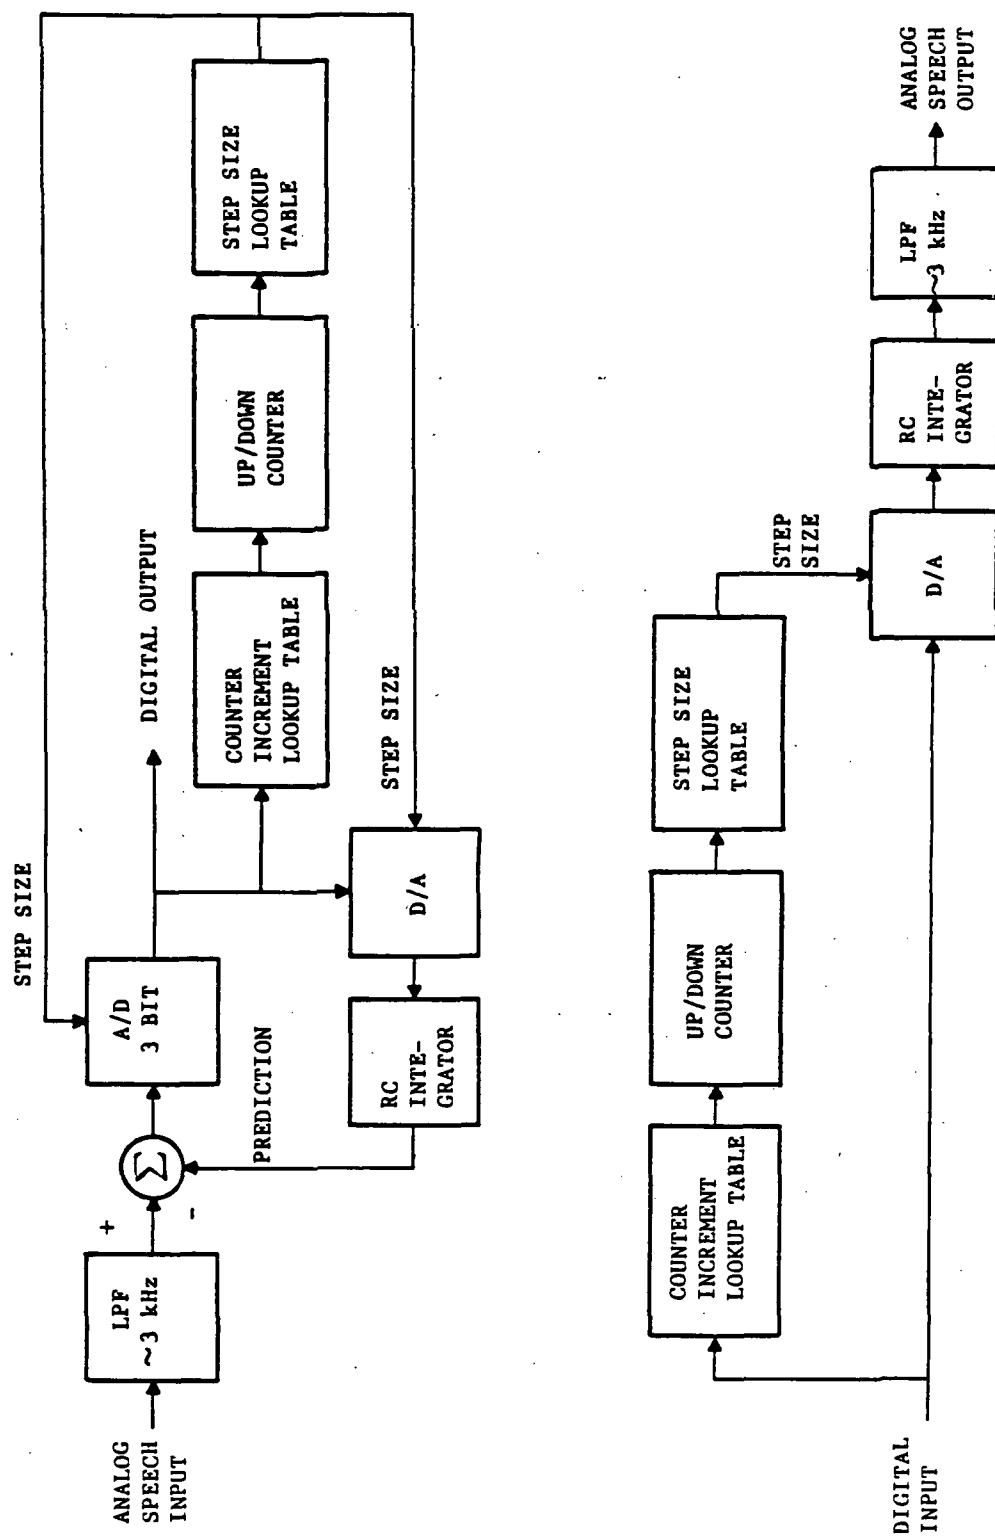


Figure 3-10. - Speech processing module.

Although admittedly a complicated structure, this buffered interface is necessitated by the relative independence of operation required between the speech processing module and the other modules of the VRU. As the sections below point out, the controller/coder, while performing simple operations, nevertheless must do so precisely at a rate tied to the fundamental speech sampling rate, lest the speech processed by the VRU take on a distorted quality. By comparison, the tasks performed by the message formatting and message entry modules are considerably more complicated but are under no such regularity constraint. Reconciliation of these disparate performance requirements leads naturally to the indicated interface structure.

3.3.2.1 Coder. The principles of waveform coding in general and adaptive differential PCM in particular are well documented in the literature (Section 3.3-1) and are not repeated here. The precise version of the technique implemented in the speech processing module, that which uses feedback adaptive quantization and fixed prediction, is shown in block diagram form in Figure 3-11. In encoding a given speech sample, it is the difference between this sample and a prediction of it which is actually encoded, hence the name differential PCM. Furthermore, the quantizer which performs the encoding has its step-size adjusted according to the magnitude of the last encoded difference, hence the adjective adaptive. The result is a speech coder which is only modestly more complex than an ordinary PCM coder but which displays equivalent performance at significantly lower bit rates. Indeed, listening tests support the use in APAS of ADPCM with three bits per sample and a sampling rate of 6.720 kHz, or a total data rate of 20.160 kilobits per second of speech.

With the exception of the low pass filters, the ADPCM coder of Figure 3-11 can, in principle, be implemented entirely in software. It is demonstrable, however, that an 8080 microprocessor operating at a clock rate of 2.1504 MHz is too slow for this task, even were it not required to perform other tasks beyond speech processing. In fact, a prudent design approach leads to implementation of the entire coder in special purpose hardware, a schematic for which is shown in Figure 3-12. This particular design uses a multiplying DAC for both analog-to-digital and digital-to-analog conversion and a ROM for implementing an adaptive quantizer via table hookup. The salient feature of the design which reduces significantly processing demands on the 8080 is the use of three 8-bit shift registers as buffer memory in communicating with the controller. These three registers hold eight speech samples, the three bits comprising each sample allocated one to a register and shifted in synchronism through the registers at the fundamental speech sampling frequency of 6.720 kHz. Control of the coder is effected through two basic signals received from the controller. The first of these control signals sets the status of the coder: active, when speech processing is to occur; and idle, otherwise. The second signal controls the mode of the coder: decode or encode. (The mode signal is ignored in idle status.)



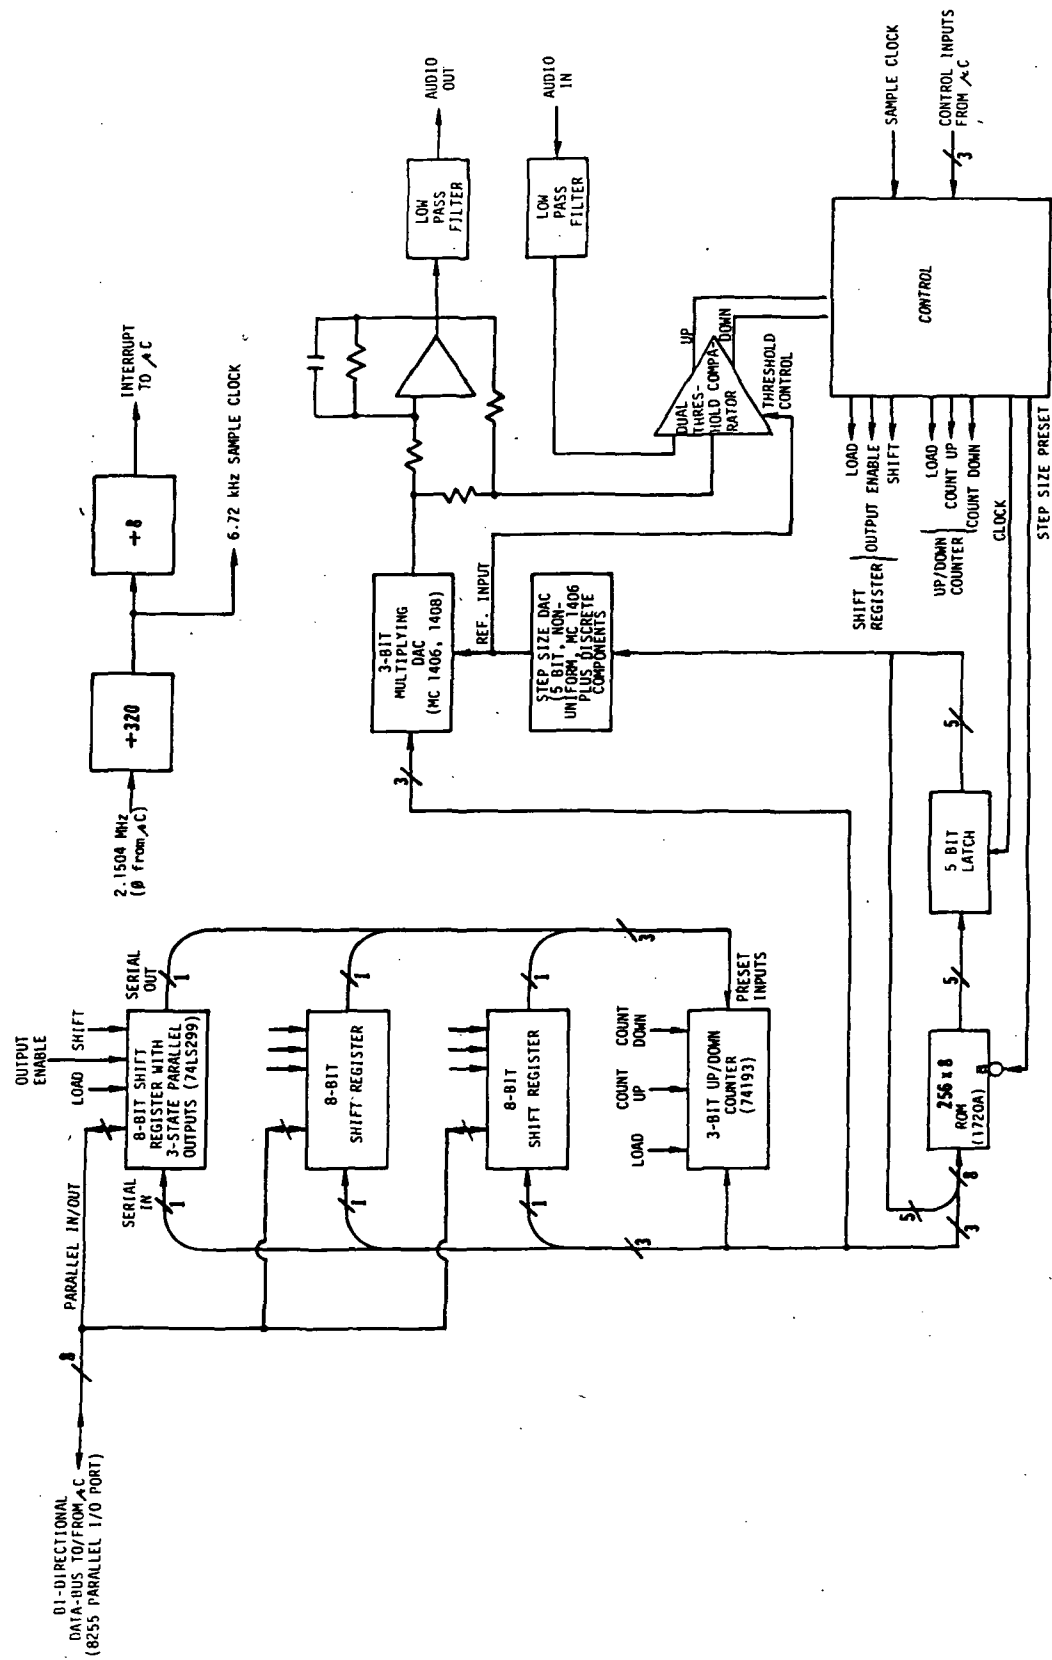


Figure 3-12. - Coder.

Timing control for the coder is obtained from the master clock operating at 2.1504 MHz. Dividing this clock by 320 provides the basic clock for the coder, operating at 6.720 kHz or precisely the speech sampling rate. Further division by 8 gives the timing control for communication between the coder and the controller. Every 1.19 msec the coder generates a signal (an interrupt) to the controller requesting a data transfer (in active status) and the appropriate control action to continue. In response, the controller performs the appropriate data transfer and issues a command to the coder which it is to follow during the next 1.10 msec.

It is noted specifically in this design approach that the coder and, hence, the controller are never quiescent; even in idle status, the coder generates a stream of interrupts to the controller to obtain directions on how to continue. In this way the controller is forced into a periodic examination of the command file for the presence of new commands with no specific wakeup action required of the data handler.

3.3.2.2 Vocabulary store. The VRU memory for all system words is a 128 kilobyte RAM capable of storing approximately 52 seconds of speech (3 bytes per 1.10 msec of speech). At the time of system startup, it is loaded with the permanent APAS vocabulary. (Many of the words of the vocabulary exist for experimental purposes, e.g., the phonetic alphabet; the conflict alert option is implemented but has never been formally tested). The length of each of the vocabulary words is indicated in the figure in terms of byte-triples. (Recall that the coder processes speech in units of three bytes.) From this information one determines that the permanent vocabulary consists of 41.9 seconds of speech and requires 105,504 bytes for its storage. Remaining, therefore, from the 131,072 bytes available in the vocabulary store are 25,568 bytes (effectively 25,566) for the storage of a discretionary message of up to 10.1 sec in length.

It should be evident that preparation of the permanent vocabulary is beyond the capabilities of the VRU, requiring as it does careful editing of individual spoken words to define their boundaries, a process best done manually using visual displays of speech waveforms. This task is therefore performed independently of the VRU, the result of which are two cassette tapes entered into the system at startup. One tape contains the ADPLM-encoder speech and is read sequentially into the vocabulary store; the other is a dictionary describing the words on the vocabulary tape as they will appear in memory according to identification number, starting address, and length (in byte-triples). Information in the VRU dictionary is used by the data handler in preparing commands for the controller.

3.3.2.3 Controller. The function of the controller, it is recalled, is basically twofold: to fetch and interpret commands from the command file and, in conjunction with the coder, to carry out these commands. It is also recalled that the controller is implemented in software as an interrupt handler and must respond to a steady stream of

interrupts from the coder occurring every 1.10 msec. The periodic interrupts create a succession of time windows, roughly 1.0 msec in duration, in each of which the controller must complete its assigned tasks and release control of the microcomputer. In order that it be able to operate on this strict schedule, the tasks of the controller are carefully segmented into a sequence of smaller subtasks, each of which can be comfortably completed in the allotted time. The result is the state transition diagram of Figure 3-13 describing the operation of the controller.

At system startup, an abort signal is posted to the controller by the VRU timing module to initialize its state. When first invoked, then, by an interrupt from the coder, the controller, recognizing the abort signal, immediately enters the reset state, sets the coder to idle status, and zeros its shift registers. In the time window following the next interrupt, the controller enters the idle state, here to remain until a new command is posted by the data handler in the command file. While in the idle state, therefore, the controller simply checks the command file for a new command, at the same time holding the coder in idle status.

Having recognized a new command, the controller proceeds through a succession of states in which it prepares to execute the commands. At the next interrupt, therefore, the controller moves to the first of these states, data acquisition, and retrieves the new command. The coder status remains unchanged in idle. The next interrupt finds the controller moving to the data interpretation state, in which the new command is broken into its constituent parts and preparation made for data transfer. At this time the coder is set to active status in the appropriate mode of operation.

In the next time window, the controller finally enters either of two data transfer states, word transfer or pause transfer, and remains in the appropriate state until the transfer is complete. In the former state, three bytes of data are passed between the vocabulary store and the coder in each time window, from the vocabulary store to the coder in decode mode and in the opposite direction in encode mode. In the pause transfer state, the shift registers of the coder are kept filled with zeros in decode mode; in encode mode, the registers are unloaded but the data ignored.* Data transfer is terminated and returned to the idle state made when a counter internal to the controller, initialized at the word/phrase length in the data acquisition state and decremented each succeeding time window, reaches zero.

Testing confirms that the controller uses sufficiently little of each time window that the microcomputer still has ample time for its other chores. In particular, the job of formatting a message can always be completed in a fraction of the time that is required to verbalize the message. The implication of this fact is that message formatting can be performed sentence by sentence as sentences are verbalized by the

*Recall that the speech processing module is never required to encode a pause.

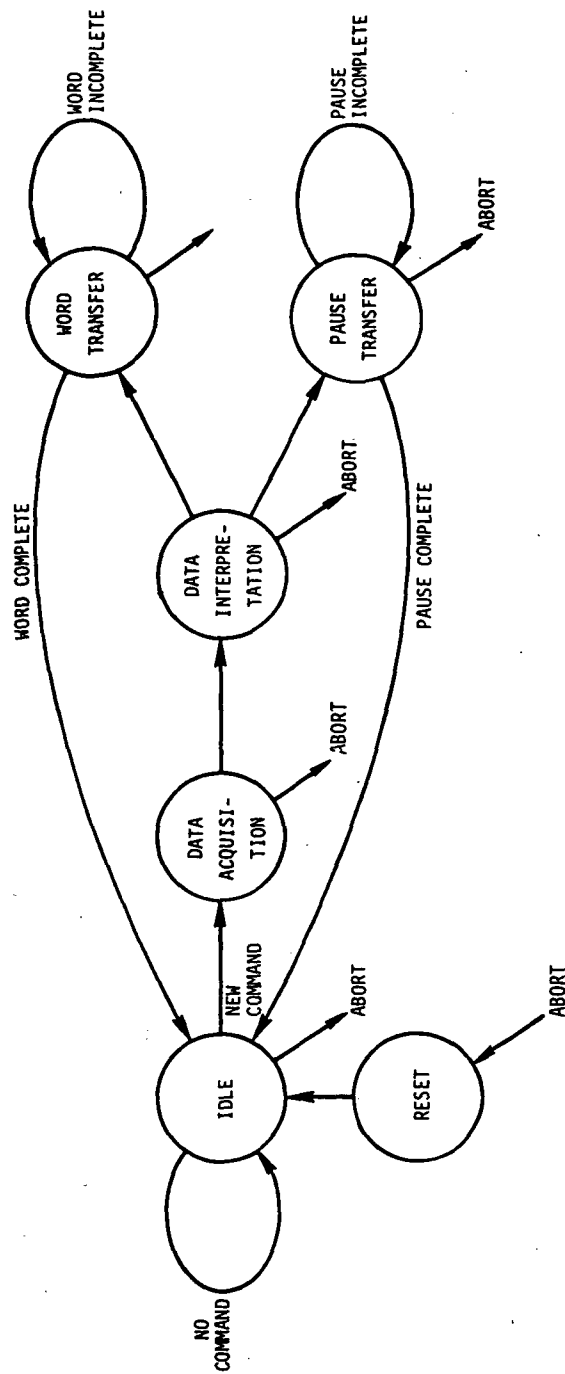


Figure 3-13. - State transition diagram.

speech processing module, guaranteeing that data used in a sentence, particularly those referring to aircraft position, are as up-to-date as possible. To exploit this situation, a sentence counter is maintained in the VRU, a counter which is incremented by the VRU timing module and decremented by the controller. The sentence counter, available always to the other modules of the VRU for inspection, allows these modules to coordinate their activities with the progress of the speech processing module.

3.3.2.4 Data handler and command file. Implemented as a subroutine in the micro-computer, the data handler is invoked by the message formatting and message entry modules with information concerning a word or pause to be played back or recorded and, in the former case, whether it is to be transmitted or not. The data handler, through consultation of the VRU dictionary, expands this information into a command to the controller and posts the result in the command file.

The commands constructed by the data handler have the logical structure shown in Figure 3-14. The two control words indicated in the figure are distinguished primarily in their destinations. The hardware control word is targeted for the coder (to set its status and mode) and the transmitter (to key transmission). The software control word, on the other hand, is used only by the controller to set its mode (playback or record) and its state (word transfer or pause transfer) and to indicate end-of-sentence. Whenever it sets the end-of-sentence bit in a software control word, the data handler also increments the sentence counter; when the command is processed by the controller, it decrements the sentence counter. At all times, therefore, this counter indicates the excess of sentences formatted over sentences verbalized.

The remaining two items in Figure 3-14 specify a starting address in the vocabulary store at which data retrieval or storage is to begin and the length of that particular word in byte-triples. The length, of course, specifies the number of consecutive time windows in which the controller is to reside in the word/pause transfer state before exiting to consult the command file for a new command. Note that the starting address is irrelevant for a pause.

3.3.3 Message Formatting Module

The message formatting module controls the content of all audio output produced by the VRU as well as the routing of this output to either the VHF transmitter or the operator control panel loudspeaker (or both). It accepts requests from the VRU timing module for the generation of any one of the three types of APAS messages, acquires the data needed for this type of message, assembles the appropriate sequence of words and control signals, and communicates the results to the data handler of the speech processing module.




Figure 3-14. - Message control list data block structure.

3.3.3.1 Message structuring. The algorithm underlying the operation of the message formatting module is an extension of the well known and intuitive fill-in-the-blanks technique. That such an approach cannot be followed directly in APAS follows from an examination of, say, the sentence concerning wind conditions in airport advisories. This particular sentence can take five distinct forms:

"WIND NOT AVAILABLE."
"WIND LIGHT AND VARIABLE."
"WIND TWO FIVE ZERO."
"WIND TWO FIVE ZERO AT ONE FIVE."
"WIND TWO FIVE ZERO AT ONE FIVE GUSTING TO TWO TWO."

The obvious problem is which of these forms one adopts to fill in the blanks. The last is the only real candidate, and yet its choice precludes the other four types of wind announcement desired in APAS.

The key to developing an appropriate technique rests on the observations that each of the sentences above can be created from six segments, namely:

[WIND]
[NOT AVAILABLE]
[LIGHT AND VARIABLE]
[TWO FIVE ZERO]
[AT ONE FIVE]
[GUSTING TO TWO TWO]

With respect to these segments, each of the sentence form is resumed in the generic form:

[WIND] [NOT AVAILABLE] [LIGHT AND VARIABLE]
[TWO FIVE ZERO] [AT ONE FIVE] [GUSTING TO TWO TWO]

consisting of the six segments in a proper order.* To create a given sentence form, one simply omits the inappropriate segments.

It is further observed that each of the six segments can be constructed using the fill-in-the-blanks techniques. For example, the segment [GUSTING TO TWO TWO] is created by filling in the blank in the form [GUSTING TO _____] with the appropriate verbalization ("TWO TWO") of wind gusts.

Careful consideration of the other sentences comprising airport and traffic advisories reveals that each can be fit into this same framework, i.e., a series

*e.g., wind speed always follows wind direction in sentences in which both appear.

of optional segments consisting of key words separated by blanks to be filled in. Virtually by default the discretionary message playback does also, consisting of one sentence of one optional segment of one word. The complete list of segmented sentence forms is given in Figure 3-15. In this figure, data are indicated as verbalized in one of several ways:

- (1) in terms of digits, e.g., hours
- (2) in forms of one of a series of words, e.g., direction or ACTIVE RUNWAY/FAVORED RUNWAY
- (3) in terms of a single word which is either included or not, e.g., NOT AVAILABLE or message

Note that the inclusion of the discretionary message in an airport advisory is data-dependent; its inclusion in a discretionary message playback is not.

Successful implementation of the sentence segmentation scheme obviously requires some basis upon which to make the omit/include decision concerning each segment comprising a generic sentence form. Since such decisions ultimately depend on data supplied by the tracking data and weather data units, a natural approach is to incorporate segment-omission indicators in this data flow. Specifically, all segments which are optional in a generic sentence form are assigned at least one piece of data. Any segment whose data are partially or totally missing (whether in a literal or figurative sense) is omitted; only segments, all of whose data are present, are retained.

For example, with respect to the segment [AT ONE FIVE] referred to above, the data in question consist only of wind speed. If the wind speed sensor is declared non-operational, then the wind speed datum is labeled missing and the segment subsequently omitted from the wind sentence. Otherwise wind speed is converted to the words "ONE FIVE" and the segment included. As an alternative example in the wind sentence, the segment [NOT AVAILABLE] has associated with it the indicator reflecting the operational status of the wind direction sensor. If the sensor is operational, then this datum is declared missing to omit the segment. If not, it is verbalized with the words "NOT AVAILABLE." (This particular technique can be used routinely for indicator variables which are not verbalized explicitly, e.g., that variable referring to pattern direction in the segment [RIGHT-HAND PATTERN].)

3.3.3.2 Data handling. A diagram showing the flow of data through the message formatting module is given in Figure 3-16. In essence, raw data from the world external to this module are first preprocessed in order to generate preformatted data consisting of the original raw data suitably modified and augmented to incorporate segment-omission indicators. From this pool, variables are retrieved as needed to fill

1.0 AIRPORT ADVISORIES

1.1 Introduction

[AIRPORT ADVISORY] [airfield]

1.2 Time

[GEE-EM-TEE hours minutes]

1.3 Wind Conditions

[WIND] [NOT AVAILABLE] [LIGHT AND VARIABLE]

[direction] [AT speed] [GUSTING TO gusts]

1.4 Altimeter Setting

[ALTIMETER] [NOT AVAILABLE] [altimeter]

1.5 Temperature

[TEMPERATURE] [NOT AVAILABLE] [temperature]

1.6 Dewpoint

[DEWPOINT] [NOT AVAILABLE] [dewpoint]

1.7 Runway Status

[ACTIVE RUNWAY/FAVORED RUNWAY number] [RIGHT-HAND PATTERN]

[CHANGING TO number] [RIGHT-HAND PATTERN]

1.8 Discretionary Message

[message]

2.0 TRAFFIC ADVISORIES

2.1 Introduction

[TRAFFIC ADVISORY] [NOT AVAILABLE]

2.2 Runway Status

[ACTIVE RUNWAY/FAVORED RUNWAY number] [RIGHT-HAND PATTERN]

[CHANGING TO number] [RIGHT-HAND PATTERN]

2.3 No Aircraft Announcement

[ZERO AIRCRAFT]

2.4 Aircraft Counts

[count AIRCRAFT ON FINAL]

[count AIRCRAFT ON BASE]

[count AIRCRAFT ON DOWNWIND]

[count AIRCRAFT ON CROSSWIND]

[count AIRCRAFT ON UPWIND]

2.5 Aircraft Description

[DEPARTING] [AIRCRAFT range MILES bearing HEADING direction]

[ABOVE PATTERN ALTITUDE]

Figure 3-15. - Segmented sentence forms.

3.0 DISCRETIONARY MESSAGE PLAYBACK

3.1 Discretionary Message

[message]

Figure 3-15. - Concluded




Figure 3-16. - Data flow.

blanks in segments, missing data triggering segment deletion. When they are used in formatting a segment, data are converted into the proper word strings, or variable phrases. The variable phrases are then merged with fixed key words to form complete segments. After each segment is formed, the individual words comprising it, along with the appropriate control signals, are communicated to the speech processing module through its data handler.

The key to a flexible message formatter--one whose operation is largely insensitive to the precise number and form of the messages it creates--lies in the manner in which input data are handled. In this connection, two of the operations just identified are critical: the preformatting operation and the conversion operation. By contrast, the other operations, e.g., phrase merging, can be made totally data independent provided preformatting and conversion are handled properly.

The preformatting operation in which raw data are converted into a form suitable for routine manipulation is a two-stage process. The first stage is generally responsible for acquiring all data required to construct a given type of message. Stated another way, this first stage involves creating a pool of raw data, on the basis of which formatting of a given type of message can proceed.

The second stage of preformatting proceeds on a sentence-by-sentence basis as a message is formatted. In this stage, data peculiar to the sentence being formatted are extracted from the message data pool formed earlier and processed as a group to create a sentence data pool. The latter data pool contains, in addition to the information from the message data pool, the missing data indicators for segment deletion. In contrast to the free form structure of the message data pool, that of the sentence data pool is based on a fixed set of storage locations, unique to the sentence in question, in which individual pieces of data come to occupy preassigned slots as indicated in Figure 3-17. This figure is largely a modification of the earlier Figure 3-15 in which all the data required to construct a given sentence of a given message are identified and catalogued.

The two stages of preformatting just described are implemented in a pair of routines, one pair per message type. The first routine of each pair, the acquisition routine, assembles the message data pool and is invoked once per message; the second routine, the processing routine, forms the sentence data pools and is invoked once per sentence. These routines are quite dependent, of course, on the data which are to appear in APAS messages, but they are affected only indirectly by the precise message formats being employed. As currently implemented, the routines are sufficiently comprehensive in the data they assemble as to be unaffected by minor format changes.

The second type of preformatting routine, the processing routine, performs two additional functions of importance to traffic advisories. The first concerns sentence form 2.5 in Figure 3-15, which is used to describe individual aircraft and may therefore appear an indefinite number of times in a traffic advisory, depending on the number of

- 1.0 AIRPORT ADVISORY
 - 1.1 Introduction
 - 1.1.1 Airfield
 - 1.2 Time
 - 1.2.1 Hours
 - 1.2.2 Minutes
 - 1.3 Wind Conditions
 - 1.3.1 Fault indicator
 - 1.3.2 Light-and-variable indicator
 - 1.3.3 Direction
 - 1.3.4 Speed
 - 1.3.5 Gusts
 - 1.4 Altimeter Setting
 - 1.4.1 Fault indicator
 - 1.4.2 Altimeter
 - 1.5 Temperature
 - 1.5.1 Fault indicator
 - 1.5.2 Temperature
 - 1.6 Dewpoint
 - 1.6.1 Fault indicator
 - 1.6.2 Dewpoint
 - 1.7 Runway Status
 - 1.7.1 Number of runway options
 - 1.7.2 Current-runway number
 - 1.7.3 Current-runway pattern direction indicator
 - 1.7.4 Next-runway number
 - 1.7.5 Next-runway pattern direction indicator
 - 1.8 Discretionary Message
 - 1.8.1 Existence indicator
- 2.0 TRAFFIC ADVISORY
 - 2.1 Introduction
 - 2.1.1 Airfield
 - 2.2 Radar Status
 - 2.2.1 Fault indicator
 - 2.3 Aircraft Counts

Figure 3-17. - Sentence data pools.

- 2.3.1 Total aircraft
- 2.3.2 Total pattern aircraft
- 2.3.3 Aircraft on final
- 2.3.4 Aircraft on base
- 2.3.5 Aircraft on downwind
- 2.3.6 Aircraft on crosswind
- 2.3.7 Aircraft on upwind
- 2.3.8 Aircraft over runway
- 2.4 Aircraft Description
 - 2.4.1 Identification
 - 2.4.2 Range
 - 2.4.3 Bearing
 - 2.4.4 Height
 - 2.4.5 Speed
 - 2.4.6 Heading
 - 2.4.7 Departure/arrival indicator
 - 2.4.8 Pattern leg
 - 2.4.9 Above-pattern altitude indicator
 - 2.4.10 Over-runway indicator
- 2.5 Runway Status
 - 2.5.1 Number of runway options
 - 2.5.2 Current-runway number
 - 2.5.3 Current-runway pattern direction indicator
 - 2.5.4 Next-runway number
 - 2.5.5 Next-runway pattern direction indicator
- 3.0 DISCRETIONARY MESSAGE PLAYBACK
 - 3.1 Discretionary Message -no data-

Figure 3-17. - Concluded

non-pattern aircraft warranting separate description. In handling such a sentence form, it is the responsibility of the processing routine to indicate when no further aircraft remain to be described, i.e., when the message data pool has been exhausted in this respect.

The second additional function, related to the first, concerns the possibility that a lengthy traffic advisory may entail the use of aircraft descriptive data which are outdated by the time they are verbalized. To combat this problem, the processing routine has the capability of updating the information on the particular aircraft being described through inspection of the files of the tracking data unit. In this respect, the routine assumes some of the duties of the acquisition routine which assembled the original data on the aircraft in question. Updating allows for the possibility that an aircraft has been dropped by the tracking data unit, in which case it is not mentioned in the traffic advisory.

The second aspect of data handling in the message formatting module, data conversion, is handled by a set of general purpose routines, so-called conversion routines, each specifically tailored to convert a numerical quantity into a sequence of words. The collection of conversion routines currently implemented in APAS is listed in Figure 3-18. For the most part, these routines compute indices of words located in the VRU vocabulary. For example, conversion routine number 14, which converts an angle in degrees into a compass bearing, computes an index between 0 and 7 specifying a category in the table:

0	North
1	Northeast
2	East
3	Southeast
4	South
5	Southwest
6	West
7	Northwest

in which a given angle falls. This index is then added to a base index--the identification number of the word "NORTH"--to arrive at the identification number of the appropriate word to verbalize the angle.

3.3.3.3 Message specification. In terms of the message formatting concepts introduced above, a given type of message is structured according to the following hierarchy of grammatical elements:

Conversion routines. The first routine is used for indicators such as those applying to [NOT AVAILABLE] or [LIGHT AND VARIABLE].

1. Do nothing (null conversion routine).
2. Select airfield name.
3. Select "ARRIVING" or "DEPARTING."
4. --NOT ASSIGNED--
5. --NOT ASSIGNED--
6. Select pattern leg.
7. --NOT ASSIGNED--
8. --NOT ASSIGNED--
9. Convert number into three digits (no sign, no zero suppression).
10. Convert number into two digits (no sign, no zero suppression).
11. Select "ACTIVE RUNWAY" or "FAVORED RUNWAY."
12. --NOT ASSIGNED--
13. Convert angle into clock position.
14. Convert angle into compass bearing.
15. --NOT ASSIGNED--
16. --NOT ASSIGNED--
17. Convert number into sign plus four digits (zero suppression).
18. Convert number into decimal form (digit, "point," digit).
19. --NOT ASSIGNED--
20. --NOT ASSIGNED--

Figure 3-18. - Conversion routines.

- sentences
 - (optional) segments
 - fixed words and variable phrases

Taken together, the last two elements comprise a fill-in-the-blanks approach to message formatting. By making the inclusion of segments optional, however, more sophisticated sentence structures are obtainable with little increase in overall formatting complexity.

In keeping with the desire to make the message formatting module flexible, the specification of messages, by which is meant a description of the elements above, is made part of input data fed to APAS at startup. The data in question divide conceptually into two parts: the first part describes vocabulary structure and the second grammatical structure. Specifically, the data appear according to the following outline:

- Vocabulary Structure
 - Dictionary
 - Conversion Identifiers
- Grammatical Structure
 - Segments
 - Sentences
 - Messages

The VRU dictionary, described earlier, is of only indirect interest to the message formatting module and is not discussed further here. The remaining items in the outline are described in the following paragraphs. (To avoid burdensome implementation details, the message specifications set forth below, although conceptually correct, deviate somewhat from those actually used in APAS).

Conversion identifiers specify to the data conversion routines described in the preceding section the location in the VRU dictionary of key words needed by the individual routines. For example, that conversion routine which translates an angle into a compass bearing requires the identification number of the word "NORTH"; this number then serves as an offset to be added to an octant number (between 0 and 7) to arrive at the proper identification number for the word describing this octant. By linking the VRU dictionary with the conversion routines in this way, changes can be made in the dictionary without impacting the conversion routines; were identifiers coded into the routines--an alternative approach--such independence would be lost.

The list of conversion identifiers is shown in Figure 3-19. The numbers shown in this figure are to be related to the conversion routines listed in Figure 3-18 and the VRU dictionary listed in Figure 3-9.

1.	0
2.	2
3.	48
4.	
5.	
6.	41
7.	
8.	
9.	22.
10.	22
11.	19
12.	
13.	23
14.	55
15.	
16.	
17.	21, 22
18.	22, 52, 22
19.	
20.	

Figure 3-19. - Conversion identifiers.

Segments are specified in terms of strings composed basically of two distinct types of elemental strings, those specifying fixed (key) words and those specifying variable (data-dependent) phrases. The former type of elemental string consists of two symbols: the letter "F" followed by the identification number of a word in the VRU dictionary. Thus "F 12" specifies the word "WIND" according to Figure 3-9. Strings specifying variable phrases, on the other hand, consist of three symbols: the letter "V" followed by a location number in the sentence data pool followed by the number of a conversion routine. Thus the string "V 3 9" will, when interpreted by the message formatting module, lead to the verbalization of the data present in location 3 of the sentence data pool by conversion routine Δ9. Reference to Figures 3-17 and 3-18 indicate that this action would be appropriate in formatting wind direction in an airport advisory. It is noted in this case that no direct reference is made to the VRU dictionary; specification of a conversion routine establishes this reference through the conversion identifiers specified for the routine.

Specification of a segment is accomplished by concatenating elemented strings, prefacing the resulting string by a segment identification number and terminating it with a stop character ("."). Thus "4 F 12." specifies that segment number 4 consists of the single fixed word "WIND," i.e., segment number 4 is [WIND]. By contrast, "7 V 3 9" specifies that segment number 7 consists of the single variable phrase described above, ostensibly referring to the segment [direction]. As a final example, "9 F 14 V 5 17." ostensibly specifies the segment [GUSTING TO gusts].

A segment specification scheme appropriate to the sentence forms of Figure 3-15 is shown in Figure 3-20. Interpretation of these specifications as indicated is a straightforward, if tedious, task performed by consulting the VRU dictionary (Figure 3-9), the structure of sentence data pools (Figure 3-17), and the conversion routines (Figure 3-18). It is faultily assumed in interpreting any specification containing a "V," i.e., one referring to a data dependent segment, that the appropriate sentence data pool is referenced. For example, specification number 7 refers to the segment [direction] only if one has in mind sentence data pool number 1.3 in Figure 3-17.

Sentences or sentence forms are specified in the natural way as sequences of segment numbers. Each such sequence is headed by a sentence identification number and terminated by a stop character. Thus, for example, "1 1 2" specifies the sentence form: [AIRPORT ADVISORY] [airfield].

Analogous to the specification of a variable phrase, however, that of a sentence must describe the data handling to accompany the formatting of that sentence. Such is indicated by three control variables immediately following the identification number and indicating, respectively, repetition type, update type, and data type. The first two of these control variables, although defined and used in a general way, practically

1.	F 252 F 1	[AIRPORT ADVISORY]
2.	V 1 2	[<u>airfield</u>]
3.	F 7 V 1 10 V 2 10	[GEE-EM-TEE <u>hours minutes</u>]
4.	F 12	[WIND]
5.	F 8 V 1 1	[NOT AVAILABLE]
6.	F 15 V 2 1	[<u>LIGHT AND VARIABLE</u>]
7.	V 3 9	[<u>direction</u>]
8.	F 13 V 4 17	[AT <u>speed</u>]
9.	F 14 V 5 17	[GUSTING TO <u>gusts</u>]
10.	F 16	[ALTIMETER]
11.	V 2 17	[<u>altimeter</u>] or [<u>temperature</u>] or [<u>dewpoint</u>]
12.	F 9	[TEMPERATURE]
13.	F 10	[DEWPOINT]
14.	V 1 11 V 2 10	[<u>ACTIVE RUNWAY/FAVORED RUNWAY</u> <u>number</u>]
15.	V 3 1	[<u>RIGHT-HAND PATTERN</u>]
16.	F 19 V 4 10	[CHANGING TO <u>number</u> <u>RIGHT-HAND PATTERN</u>]
17.	V 5 1	[<u>RIGHT-HAND PATTERN</u>]
18.	F 251 V 1 1	[<u>message</u>]
19.	F 252 F 38	[TRAFFIC ADVISORY]
20.	V 1 17 F 39	[<u>ZERO AIRCRAFT</u>]
21.	V 3 17 F 39 F 41	[<u>count</u> AIRCRAFT ON FINAL]
22.	V 4 17 F 39 F 42	[<u>count</u> AIRCRAFT ON BASE]
23.	V 5 17 F 39 F 43	[<u>count</u> AIRCRAFT ON DOWNWIND]
24.	V 6 17 F 39 F 44	[<u>count</u> AIRCRAFT ON CROSSWIND]
25.	V 7 17 F 39 F 45	[<u>count</u> AIRCRAFT ON UPWIND]
26.	V 7 1	[DEPARTING]
27.	F 39 V 2 18 F 51 V 3 14	[AIRCRAFT <u>range</u> MILES <u>bearing</u> <u>HEADING direction</u>]
28.	F 46 V 9 1	[<u>ABOVE PATTERN ALTITUDE</u>]
29.	F 251	[<u>message</u>]

Figure 3-20. - Segment specification scheme.

refer only to the aircraft description sentence in traffic advisories. As their names indicate, the former variable indicates whether or not the sentence is used repetitively; the latter, whether or not data which will appear in the sentence data pool is to be updated. Both variables take one of two values: "Y" signifying repetition or updating or "N" for the opposite cases. The third control variable, data type, specifies the set of data to be assembled in the sentence data pool.

The previous example, therefore, should be expanded to read: "1 N N 1 1 2." signifying no repetition, no updating, and data set number 1. The complete set of sentence specifications is given in Figure 3-21.

Messages are specified simply in terms of sentences. The list of message specifications is given in Figure 3-22.

It should be noted that there exists considerable latitude within the basic structure of the message formatting module for accomplishing the same objectives in several different ways. Thus the message formatting rules just set forth are not unique: given the arrangement of data within the sentence data pool shown in Figure 3-17 and the conversion routines of Figure 3-18, there still exist many alternatives for specifying the desired APAS message set.

To a great extent, it is this inherent latitude which allows changes in message formats to be effected through input data alone.

3.3.4 Message Entry Module

Conceptually, the message entry module parallels the message formatting module and does for message input what the latter does for message output. Practically, however, the message entry module is quite limited in its operation: it formats the input of only a single word in the VRU vocabulary, namely the discretionary message. Thus the control information it communicates to the data handler of the speech processing module consists of a single word (number 251) to be encoded.

A subsidiary activity performed by the message entry module consists of using the system operator to begin speaking. Thus this module is also responsible for operation of the operator control panel lights according to procedures described earlier (Section 2.2.2).

3.3.5 VRU Timing Module

As the scheduler of all input and output messages of the VRU, the VRU timing module can be viewed as the principal controller of APAS. It is implemented as the main program of the microcomputer, invoking, as necessary, the subroutines which realize the functions of the message formatting and message entry modules.

1.	N N 1 1 2	1.1
2.	N N 2 3	1.2
3.	N N 3 4 5 6 7 8 9	1.3
4.	N N 4 10 5 11	1.4
5.	N N 5 12 5 11	1.5
6.	N N 6 13 5 11	1.6
7.	N N 7 14 15 16 17	1.7
8.	N N 8 18	1.8
9.	N N 1 19	2.1
10.	N N 2 5	2.2
11.	N N 5 14 15 16 17	2.3
12.	N N 3 20	2.4
13.	N N 3 21 22 23 24 25	2.5
14.	Y Y 4 26 27 28	2.6
15.	N N 1 29	3.1

Figure 3-21. - Sentence specifications.

1.	1 2 3 4 5 6 7 8	(AIRPORT ADVISORIES)
2.	9 10 11 12 13 14	(TRAFFIC ADVISORIES)
3.	15	(DISCRETION MESSAGE PLAYBACKS)

Figure 3-22. - Message specifications.

The module is equipped with two timers which indicate respectively the elapsed time since the beginning of the last airport advisory broadcast or the last traffic advisory broadcast. (The timers are implemented in a single interrupt handler keyed to the system clock.) When either of these timers indicates that the elapsed time--since the last of its advisories was broadcast--has reached the nominal desired time interval between such advisories, the message formatting module is invoked to broadcast the next one of that type of advisory.

In the ongoing process, the VRU timing module continually monitors the operator control panel to respond to a request from the system operator to enter a discretionary message. If so, then as soon as any advisory broadcast currently in progress is completed, this module first calls on the message entry module to accept the discretionary message and then the message formatting module to play back the newly entered message locally for operator approval. This process is repeated until the required approval is obtained, at which time an immediate airport advisory is scheduled.

The operation of the VRU timing module is detailed more precisely in the flow diagram of Figure 3-23. Each advisory timer is seen to be initialized at the beginning of a broadcast to the nominal time interval desired between such advisories, from which point it counts down to and holds at zero. Note that the order in which these timers are checked gives precedence to airport advisories, an approach which is entirely consistent with the APAS broadcast schedules discussed earlier: either no traffic advisories are broadcast or traffic advisories are interspersed between airport advisories. Note also that traffic advisories are always timed (for the sake of uniformity) but that a secondary check is performed before their release in case they are not being broadcast.

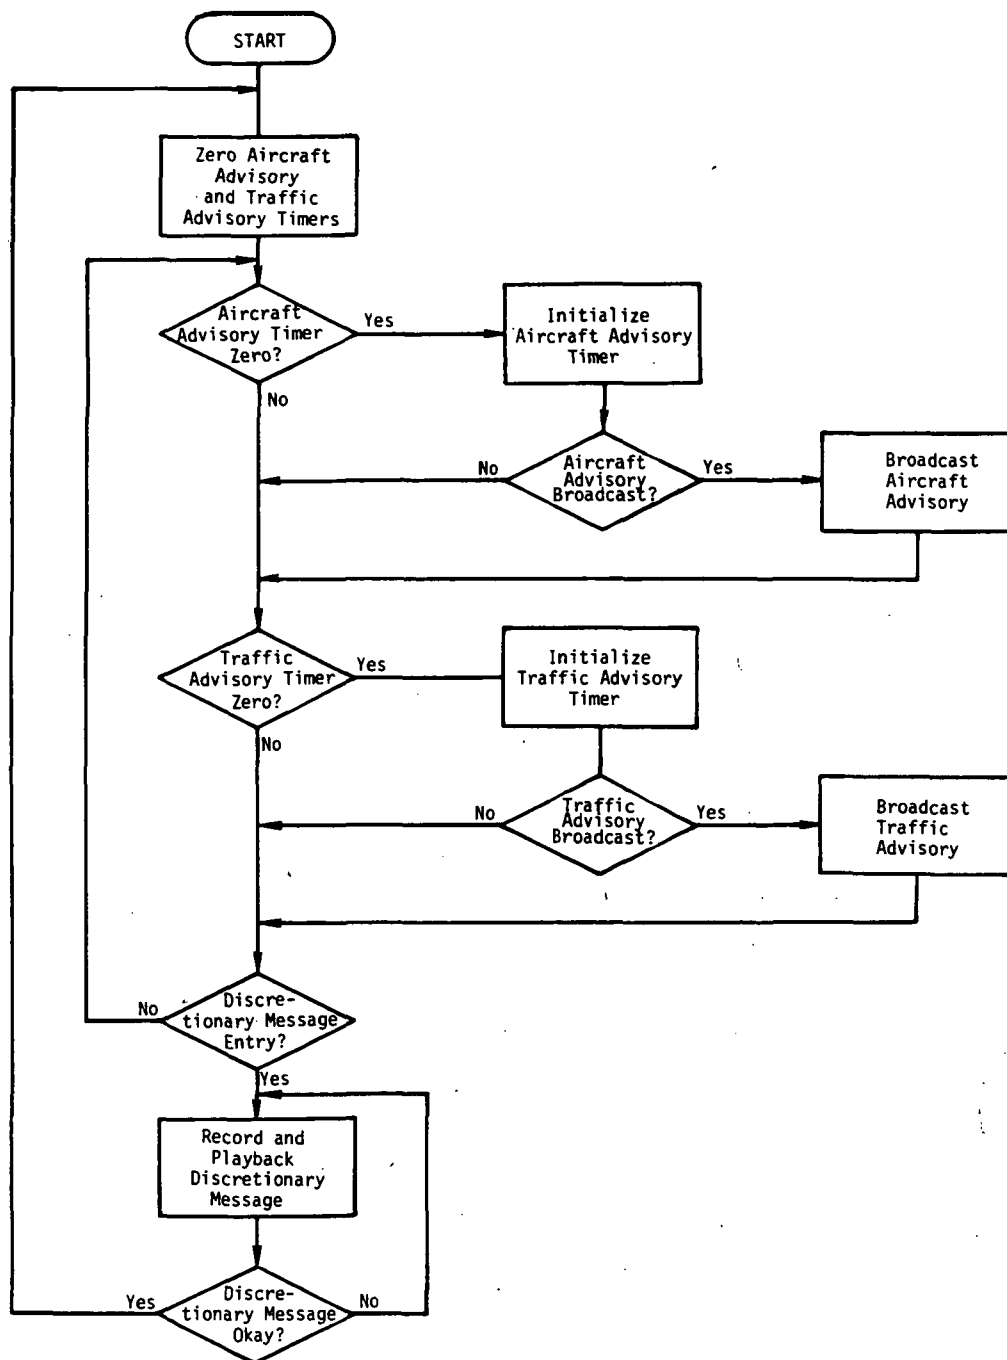


Figure 3-23. - Scheduling of advisory broadcasts.

4.0 TEST AND EVALUATION OF THE EXPERIMENTAL APAS

Test and evaluation of the experimental APAS were accomplished by a two phase program which commenced in April 1979 following integration of the radar system and the Research Triangle Institute (RTI) supplied computer and software. The initial APAS configuration was predicated upon a theoretical model, but variability had been built into both hardware and software so that various system configurations could be tested.

The objective of the first phase, low density traffic tests, was to define the optimum system configuration to meet the APAS requirements. These tests were conducted at Wallops Flight Center for one year commencing in April 1979 using three general aviation aircraft: an Aero Comanche, a Cessna 150, and a Cessna 172.

In late May 1980, the APAS was moved to Manassas Municipal Airport, Manassas, Virginia, to perform the second phase of the testing program. The objectives of this phase were to evaluate the APAS performance in a high density uncontrolled operational environment and to obtain a pilot's evaluation of the APAS. The Manassas airpark was selected because it is a high density uncontrolled airport with an estimated 200,000 operations per year. The testing at Manassas occurred between June and September 1980 with a six week continuous period in which APAS was operated eight hours each day between 0900 and 1700 hours. Following the six week period, the APAS was operated between 0900 and 2300 hours for one week to obtain data at night.

4.1 Overview

Prior to discussing the system configuration changes and experimentation, an overview of the major areas of investigation for the APAS follows:

4.1.1 Clutter

Cost considerations forced the utilization of a skin tracking surveillance type radar because the cost savings incurred in this type versus other type radars were more than the total APAS system cost. The disadvantages in using this type radar are in its target detection and tracking ability, which requires both to be done in a clutter environment. Target detection depends on receiving radar signals reflected off the target of interest, in our case, aircraft. This type radar also receives reflected signals from other targets (hereafter called interference) which are not of interest in an APAS. Sources of interference include trees, buildings, cars, weather, external radiation, etc. The effects produced by interference are two: (1) false target reports, and when signals from interference are larger than ones from aircraft, and (2) a non-detectable target. The means of combating these problems in the APAS were to utilize hardware to

minimize the interference signals and to provide detection and tracking software to operate in this environment.

4.1.2 Timing

Optimizing system timing for both traffic and airport advisories was a major area of investigation. Very early in the program it was decided that the radar antenna rate would have to be the system clock. The revolutions per minute (RPM) of the antenna established the radar data rate (Section 3.1) and therefore the target update rate. Since an aircraft's true position would be changing between target updates, a minimum update time was desired. If the time between updates got too large, the following undesirable effects would occur:

- (1) Pilots could not identify themselves in traffic reports.
- (2) False traffic reports occurred because the tracking software could not correctly associate target reports with computer tracks of aircraft. (This effect would predominately occur in the pattern where aircraft were turning from one pattern leg to the next.)
- (3) Traffic reports which made no sense, i.e., a report of an aircraft on final where the previous report was on crosswind, etc.

The target update rate was a function of both the radar antenna's RPM and the number of receive antennas used. The trade-offs in this area were:

- (1) Increasing the number of antennas extends the radar coverage area and therefore enhances the system trackability, but it also increases the time between target updates. (The APAS was designed to process one receive antenna at a time.)
- (2) Decreasing the number of receive antenna by increasing antenna beam widths would increase the clutter interference and would produce tracking problems.
- (3) Decreasing the number of receiver antennas by eliminating coverage would reverse the effect discussed in (1) above.

4.1.3 Coverage

Optimizing the coverage area was a major area of investigation. Previous studies had indicated that the pattern area extended for approximately two nautical miles about an airport, and aircraft density significantly decreased beyond this range. The experimental APAS was designed to process radar information of 64 range cells and up to 256 azimuth sectors as selected by the digitizer. The issues in this area of investigation were to determine the range cell size and azimuth subcell processing to optimize the coverage, detection, and tracking objectives of APAS. (See Reference 1).

4.1.4 Pilot Acceptance

The primary objective of the traffic advisory system is to enhance the see-and-be-seen concept. To accomplish this, the APAS requires pilots to disseminate traffic report information so that they can identify themselves and potential conflicting aircraft. Since traffic reports would be issued on all aircraft within the coverage range, pilots would have to mentally filter out those reports which would not conflict with their aircraft (i.e., it would be impossible to keep track of all traffic reports). Additionally, the pilot would be required to track the potentially conflicting aircraft from successive traffic reports so that he could identify an occasional false report. The ability of pilots to accomplish these tasks and his acceptance of the APAS concept was a major investigation area.

4.2 Low Density Traffic Test

The first phase of APAS testing started with an evaluation of the initial hardware and software configuration and ended with the configuration defined in Sections 2 and 3 of this report. The name, low density traffic test, is derived from the fact that a limited number (usually three or less) aircraft were used in controlled tests to define by experimentation the optimum APAS configuration. The specific tests performed are described in the preceding sections of this report.

4.2.1 Data Rate Tests

Data rate tests were those tests designed to determine the maximum RPM of the radar antenna. The factors which limit the RPM are as follows:

- (1) Sufficient time for the track-while-scan computer to process the radar information, i.e., perform all processing from data input to traffic definition.
- (2) Sufficient time for the digitizer to perform the following operations between the eight integrated radar pulse in an azimuth sector and the first transmitted pulse of the next azimuth sector.
 - (a) perform error checking.
 - (b) perform data comparison or addition.
 - (c) transfer data from cell A and cell B memory to cell C (output) memory.
 - (d) transfer data from C memory into track-while-scan computer.

Tests indicated that the track-while-scan computer could operate at speeds up to 33 RPM, but under design maximum traffic densities the system would saturate between 30 and 30.5 RPM. Rate tests on the digitizer indicated the system could operate in a calm wind at

speeds up to 33 RPM, but winds could produce errors at speeds in excess of 30.1 RPM. A 30.0 RPM was selected for the antenna speed.

4.2.2 Radar System Tests

The APAS radar coverage area was initially conceived analytically to survey 360° of azimuth, a range sufficient to cover the pattern, and elevation coverage sufficient to track aircraft who were within pattern range and whose altitude was less than 3,000 feet. An additional requirement was to minimize the size of range cells and azimuth sectors so that tracking accuracy would be enhanced. This coverage area was achieved by 64 range cells of 75m length (3sm total range), 256 azimuth sectors of 1.4° , and 5 receive antennas set at elevation angles 13° apart.

4.2.2.1 Azimuth tests. Timing tests on the track-while-scan computer indicated that it was marginal on whether a 8,192 words/sec. rate, produced by the 1.4° azimuth sector, could be processed in conjunction with the other processing duties of this computer. These tests led to the inclusion of a preprocessor in the digitizer which effectively halved the processing rate.

Tests were performed to determine the optimum mode, comparative or additive, for the digitizer's preprocessor. Test results indicated that the comparative mode increased the probability of target detection but, at ranges where signal returns were weak, the false alarm rate would be higher for this mode than those produced by the additive mode. False target reports at the extended ranges were easily detected by the tracker, and few false traffic reports would result. The results of these tests indicated that the comparative mode of operation was more suited than the additive one.

4.2.2.2 Range tests. Range tests were performed to determine the optimum range coverage for each of the receive antennas. For elevation antenna number 2 through the highest antenna, these tests consisted of determining the number of range cells necessary to achieve the desired altitude coverage as the number and beam width of antennas were varied.

For elevation antenna number 1, the lowest, the test consisted of determining the coverage range so that the length of each range cell could be determined. Very early in the range testing program it became obvious that factors other than analytical solution would determine the range coverage. Some of these factors were:

- (1) The coverage range must include both the pattern and a track initiation range.
- (2) The track initiation range was proportional to the track update rate and the speed and direction of arriving aircraft.
- (3) The coverage range would define the size of the range cell.
- (4) Tracking range uncertainty increases as range cell size increases.

- (5) Computer speed and memory capacity dictated that the number of range cells processed be kept to a minimum and any limitations on the number of cells processed would have to occur in the upper antenna.
- (6) The number of range cells processed in the upper antenna should consider altitude coverage and a track initiation probability factor.

The range coverage parameters which optimized the three antenna 30 RPM final APAS configuration were:

Antenna Number	Cell Size	Total Range	No. Cells Processed
1	150m	6.0sm	64
2	150m	4.0sm	42
3	150m	2.5sm	20

4.2.2.3 Antenna tests. Antenna tests were performed to determine the antenna configuration, i.e., the number of receive antennas, the elevation, and the beam widths of each. Initial testing demonstrated that the antenna configurations would depend more on clutter suppression and tracking requirements than on coverage requirements. Factors which were determined through testing include:

- (1) The number of receive antennas are limited by the target update rate when the antennas are processed serially. Detrimental tracking effects occur when the target update rate exceeds six seconds; therefore, with a 30 RPM antenna rate, three receive antennas were the maximum.
- (2) Elevation and beam widths became trade-offs between tracking, clutter suppression, and coverage.
- (3) The antenna with the lowest elevation angle should be a narrow beamwidth ($\approx 3^\circ$ to 5°) high gain one. Target reports from this antenna will occur close in range (aircraft departing the runway or on final) or at maximum range (arriving aircraft entering the radar coverage area). Therefore, the elevation angle should be as low as possible conducive with clutter suppression objectives.
- (4) The second and third antenna should be set so that their elevation angle minus their beam width is coincident with the next lower antenna's elevation angle plus its beamwidth, i.e., no gaps between antenna coverage.
- (5) The beamwidth of the second and third antenna should be as wide as possible conducive with clutter suppression requirements.
- (6) Antenna select control should be exercised by the tracking computer.
- (7) The hold at higher elevations should be handled by the tracking computer.

- (8) If parallel processing of multiple beams is employed, a fourth antenna should be used to detect maneuvering aircraft at high elevation angles. The tracking computer should only process data from this antenna when it predicts the existence of a target within its beamwidth.

4.3 High Density Traffic Tests

In late May 1980 the APAS was moved from WFC to the Manassas Municipal Airport, Manassas, Virginia, to test the system in a high density uncontrolled environment. Manassas was selected because it is a high density uncontrolled airport with an estimated 200,000 operations per year. The purpose of the Manassas testing was to determine the adequacy of system specifications and to ascertain whether any system degradation would occur due to high traffic density or other factors. The primary areas of concern were system cycle time, target detection, tracking, and message rates.

From June 23, 1980, to August 16, 1980, the experimental APAS was operated daily between 0900 and 1700 (0900 to 2300 the week of August 11). The methods employed to evaluate APAS performance included a continued verification of advisory reports and the maintenance of a system anomaly and pertinent data log. Additionally, during two 90-minute periods each day, all traffic advisory reports were recorded, and a count was obtained of those reports verified or unverified by radar or visual spotters.

The daily traffic density throughout the test period is included in the Appendix. The maximum traffic density during this period occurred on Sunday, July 13, 1980. The total track rate, operational rate, and traffic report histogram data for this day are presented in Figure 4-1 through 4-3, respectively. (Total track rate is the number of APAS validated tracks per hour; the operational rate is the sum of take-offs and landings per hour; the traffic report histogram depicts the number of traffic reports containing "N" number of aircraft). This data indicates that the APAS operated for five hours at an operational rate exceeding 50 operations per hour with a peak rate of 70 operations during a one-hour period.

4.4 Synopsis of Test Results

An evaluation of the APAS performance during the peak traffic density occurrence of July 13, 1980, indicates:

- (1) the system successfully handled its design limit of ten aircraft per traffic report on several occasions,
- (2) the 30 RPM antenna cycle time was maintained,
- (3) no degradation occurred in traffic report accuracy (the highest accuracy rate achieved during the six week test occurred during the five hour high density period on July 13, and

TRACK RATE
JULY 13, 1980




Figure 4-1. - Track rate.

OPERATIONAL RATE

JULY 13, 1980




Figure 4-2. - Operational rate.

TRAFFIC REPORT HISTOGRAM

JULY 13, 1980




Figure 4-3. - Traffic report histogram.

- (4) the time for a traffic advisory message exceeded the allotted 20-second period, but system software successfully handled this situation by delaying the next advisory by the time overrun.

The APAS performance in marginal VFR conditions was mixed. On two occasions, during very hazy conditions, the APAS experienced no performance degradation; on other occasions, in light to moderate rain, the traffic advisory system was turned off because of numerous false target reports. The APAS contains computer software which detects the existence of rain and attempts to maintain pattern reports while deleting traffic reports outside the pattern in the area where the rain occurs. This software was used with favorable results on several occasions during isolated thunderstorms. Although the computer software in the experimental APAS did not contain the proper messages, it appears that the rain detection software could be expanded to handle the moderate rain problems.

The experimental APAS had a seven-to-eighteen second system delay which resulted in aircraft completing a pattern leg turn being reported on the previous pattern leg. This time delay was caused by a combination of the traffic advisory reporting time, the six-second target update rate, and target coast mode following a missed detection. Initial users of APAS expressed concern about the delay, but pilots who continually used the system indicated that, if they didn't locate the traffic reported in a pattern leg, they would instinctively look for traffic on the next pattern leg and, therefore, the delay wasn't a problem.

During the APAS operational period, the Manassas UNICOM voice traffic was significantly reduced. This condition was illustrated by a comparison between the voice traffic which occurred immediately before, to that which occurred during short periods in which APAS messages were terminated to store tracking data. During these periods pilots used the self-announcement system. Although measurements were not made to quantize it, the reduction was significant enough to make it obvious to those who monitored the UNICOM frequency.

The only APAS anomaly occurred in the runway selection algorithm, which caused a runway change three times over a five minute period in light and variable winds. An analysis of the problem indicated that the number of runways impacted several input numbers in unforeseen ways. An immediate fix to the problem was implemented by changing the value of an input number, but this fix would negate the universality of the algorithm. A solution to this problem has been proposed but has not been tested.

Throughout the six-week test period, 95 percent of the APAS reports were verified during the 90-minute counts. The breakdown on the five percent incorrect reports showed that one percent was loss of track on the final leg, one percent was late reports on departing aircraft, two percent were false tracks caused by large

earth-moving equipment being used to construct a parallel runway, and one percent was for various other causes. The occurrences of the incorrect final and departure reports were enhanced by earth-moving equipment and site location problems unique to the experimental APAS. These two factors caused a higher-than-normal radar signal to be required for target detection, therefore, decreasing the probability of detection. It should be noted that the APAS software contains a computer code to eliminate problems produced by roadways, but it could not be utilized because the roadway for the earth-moving equipment was one-half mile wide.

4.5 Pilot Evaluation

A primary objective of the Manassas testing was to obtain pilot evaluations of the APAS concept in the uncontrolled high density environment. To accomplish this, the experimental APAS was operated for an eight-hour period each day for six weeks. An informational package, including a questionnaire, was distributed to pilots who used the system, and one hundred pilots responded to the questions (Q). Their responses (R) and an authors comment (C) are presented:

Q: Date and time of experience?

R: Not applicable

Q: Pilot Hours?

R: 50 - 5%

100 - 6%

200 - 12%

500 - 18%

1000 - 17%

>1000 - 42%

Q: a. Function?

R. Pilot - 99%

Co-Pilot - 1%

b. Rating?

Private - 7%

Commercial - 2%

Instrument - 12%

SEL - 28%

Multiple - 51%

Q: Type of aircraft?

R: SEL - 81%

MEL - 16%

Other - 3%

Q: APAS Voice Quality?

R: Unusable - 0%

Confusing - 1%

Satisfactory- 39%

Excellent - 53%

Other - 7% (4% favorable and 3% unfavorable)

Q: Was the airport advisory two minute rate satisfactory?

R: Yes - 89%

No - 11%

C: Most of the no responses occurred on hazy days when pilots indicated they needed favored runway information more often. The two-minute rate was insufficient because pilots were released from a controlled condition to VFR and tuned to the APAS broadcast after they had the airport in sight. Invariably, some pilots had to fly around the airport for almost two minutes to learn the favored runway from the next airport advisory.

Q: Was the airport advisory message format acceptable?

R: Yes - 92%

No - 8%

Q: Any improvements in airport advisory?

R: No improvement - 38%

Repeat runway more often - 12%

Runway change confusing - 10%

Temperature and dewpoint
information not necessary - 6%

Other - 34%

Q: a. Did you experience a change in active runway?

R: Yes - 18%

No - 82%

C: The APAS selects the favored runway by a technique which is a function of the prevailing winds. When conditions occur which produce a change in the favored runway, the APAS initiates the change by announcing it on the next airport advisory message. On each of the next six traffic advisory reports, which occur between airport advisories, the runway change is announced following the traffic report. The process is completed on the next airport advisory when the favored runway is announced to be the new one.

b. If so, describe your reaction.

Dangerous - 22%

Confusing - 28%

Satisfactory - 28%

Orderly - 22%

C: Two occurrences contributed negative responses to this question. The first was the runway change anomaly described in the system performance evaluation where several aircraft were forced to taxi back and forth on the taxiway, while the APAS kept changing the favored runway. This occurrence caused several responses that the runway change method was confusing.

The second occurrence resulted from a breakdown in control over the favored runway. Since controlling the runway would be part of any APAS evaluation, an agreement was made with the Manassas airport authorities, whereby the Manassas FBO would direct anyone requesting the favored runway to obtain the information from APAS broadcast. On two occasions this procedure failed and a favored runway, different than the one selected by APAS, was announced on the UNICOM frequency. On both occasions, the result produced was two aircraft simultaneously attempting to land on opposite runways. Announcements were made to divert the aircraft, but several dangerous responses were received from pilots.

Q: Was the traffic advisory rate satisfactory?

R: Yes - 89%

No - 11%

C: A non-limiting method was chosen to announce traffic information for the APAS. Non-pattern reports were ordered by azimuth so that pilots could differentiate potential conflicting and non-conflicting aircraft. This method would produce numerous target reports in high traffic densities so the next several questions were designed to evaluate the method.

Q: a. Were you able to identify yourself in the traffic advisory?

R: Yes - 95%

No - 5%

b. How many other aircraft were being reported?

1 - 9%

2 - 13%

3 - 24%

4 - 19%

5 - 19%

6 - 10%

7 - 4%

8 - 1%

Q: Were you able to locate all other traffic in the advisory?

R: Yes - 46%

No - 54%

If no, were you able to locate all traffic presenting a potential conflict?

Yes - 86%

No - 14%

Q: What is your opinion of the traffic advisory?

R: Disastrous - 3%

Confusing - 8%

Satisfactory - 34%

Wonderful - 30%

Other - 25% (19% favorable and 6% unfavorable)

Q: Did you experience any false target reports?

R: Yes - 14%

No - 86%

If yes, was it a problem?

Yes - 45%

No - 55%

Q: Did you site any traffic that was not reported by the system?

R: Yes - 20%

No - 80%

Q: Was the traffic advisory information in a format that you fully understood?

R: Yes - 95%

No - 5%

Q: What is your opinion of the APAS message vs. self-announcement?

R: Favored APAS - 87.5%

Favored self-announcement- 12.5%

Q: Comments:

R: Favorable - 86.5%

Unfavorable - 13.5%

C: The favorable comments indicated that pilots thought that APAS was a safer system than the self-announcement procedure. The unfavorable comments were in two general areas: system delay, and lack of knowledge about pilot intentions.

5.0 PROPOSED DEVELOPMENT

The testing at Manassas was the first attempt to evaluate an APAS in a high density uncontrolled environment. As a minimum, this test proved that low-cost automated systems can provide airport and air traffic advisory information at high density uncontrolled airports, and a large majority of the users preferred the APAS over a self-announcement procedure.

An evaluation of the operational performance of the APAS during the Manassas testing indicated that additional enhancements could be obtained in the following areas:

Clutter suppression. Enhancements in clutter suppression will decrease the false target report rate and could solve the final and departing aircraft reporting problem. The enhancements could be made in several ways, such as increasing the height of the antenna platform and optimizing the transmit and receive antenna elevation beam-width. It is recognized that an MTI type of radar would solve the clutter problem, and this type radar may be required at some trouble airports, but the cost of this solution should be analyzed vs. system affordability.

System delay. Decreasing the system time delay appears feasible without significantly increasing system cost by using a dual receiver radar system and concurrently processing two receive antennas. It is recommended that the lowest elevation antenna be processed every cycle and the two upper elevation antennas be alternately processed. This method should result in a three-to-seven second system delay and have additional benefits such as increasing the range of initial target reporting and decreasing the false target report frequency.

Channel assignments. The decrease in UNICOM voice traffic during APAS operations, and the APAS requirement of only a 10-to-20 nautical mile broadcast coverage area, are significant factors in accessing frequency channel assignments for an operational system. Additional channels for the APAS broadcast may be obtained by assigning more uncontrolled airports the same UNICOM frequency.

It is recommended, therefore, that a phase II program be initiated to incorporate the aforementioned enhancements into an APAS. The inclusion of these system modifications should produce an APAS which will meet the primary objective of an automated system for enhancing the see-and-be-seen rule of flight.

6.0 CONCLUSION

The APAS project was conceived to develop an automated system to enhance the see-and-be-seen rule of visual flight and which would be affordable to the thousands of municipal and privately owned airfields in the United States. In pursuance of these objectives, an experimental system was developed, tested, and demonstrated. The information derived from these processes has been described in this report, and it is hoped that it will lead to a system which will increase safety to the general aviation pilot. In concluding this report, no better conclusion can be stated than that expressed by a pilot as he witnessed the APAS performance while watching the traffic flow at Manassas: "This thing really works."

APPENDIX Traffic Density

TRACKING DATA FOR: 6/23/80

ANALYSIS TYPE

ARRIVAL

DEPARTURE

TOTAL TRACKS

TARGET HISTOGRAM

NO. AIRCRAFT REPORTS

FREQUENCY

PERCENTAGE

TOTAL TRACKS

RATE/HR.

START TIME

STOP TIME

TRAFFIC RATE HISTOGRAM

TIME OF DAY (HRS.)

ARRIVALS/HR.

DEPARTURES/HR.

TOTAL TRACKS/HR.

TRACKING DATA FOR: 6/24/80

ANALYSIS TYPE

ARRIVAL

DEPARTURE

TOTAL TRACKS

TARGET HISTOGRAM

NO. AIRCRAFT REPORTS

FREQUENCY

PERCENTAGE

TOTAL TRACKS

RATE/HR.

START TIME

STOP TIME

TRAFFIC RATE HISTOGRAM

TIME OF DAY (HRS.)

ARRIVALS/HR.

DEPARTURES/HR.

TOTAL TRACKS/HR.

TRACKING DATA FOR: 6/25/80

ANALYSIS TYPE		TOTAL TRACKS										RATE/HR.			START TIME			STOP TIME							
ARRIVAL		91										12.			0851 8.0			1704 1.1							
DEPARTURE		106										14.			0851 8.0			1704 1.1							
TOTAL TRACKS		231										30.			0851 8.0			1704 1.1							
TARGET HISTOGRAM																									
NO. AIRCRAFT REPORTS		0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20			
FREQUENCY		308	436	354	176	78	18	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0			
PERCENTAGE		22	31	25	12	5	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0			
TRAFFIC RATE HISTOGRAM																									
TIME OF DAY (HRS.)		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24
ARRIVALS/HR.		0	0	0	0	0	0	0	0	0	12	12	21	7	18	3	9	10	0	0	0	0	0	0	0
DEPARTURES/HR.		0	0	0	0	0	0	0	0	0	17	13	19	14	15	15	9	10	0	0	0	0	0	0	0
TOTAL TRACKS/HR.		0	0	0	0	0	0	0	0	0	34	32	34	15	37	20	29	29	0	0	0	0	0	0	0

TRACKING DATA FOR: 6/26/80

ANALYSIS TYPE	TOTAL TRACKS												RATE/HR.		START TIME		STOP TIME							
ARRIVAL	124												17.		0903 9.0		1635 39.6							
DEPARTURE	123												17.		0903 9.0		1635 39.6							
TOTAL TRACKS	282												38.		0903 9.0		1635 39.6							
TARGET HISTOGRAM																								
NO. AIRCRAFT REPORTS	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20			
FREQUENCY	184	429	430	191	78	12	7	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
PERCENTAGE	13	13	32	14	5	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
TRAFFIC RATE HISTOGRAM																								
TIME OF DAY (HRS.)	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24
ARRIVALS/HR.	0	0	0	0	0	0	0	0	0	21	15	26	6	26	7	17	11	0	0	0	0	0	0	0
DEPARTURES/HR.	0	0	0	0	0	0	0	0	0	14	18	25	11	23	9	19	6	0	0	0	0	0	0	0
TOTAL TRACKS/HR.	0	0	0	0	0	0	0	0	0	44	40	52	23	38	19	49	35	0	0	0	0	0	0	0

TRACKING DATA FOR: 6/27/80

ANALYSIS TYPE

ARRIVAL

DEPARTURE

TOTAL TRACKS

TARGET HISTOGRAM

NO. AIRCRAFT REPORTS 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20
FREQUENCY 140 183 175 143 77 10 2 1 2 0 0 0 0 0 0 0 0 0 0 0 0
PERCENTAGE 19 24 23 19 10 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0

TOTAL TRACKS

RATE/HR.

START TIME

STOP TIME

16.

13.

40.

0852 8.0

0852 8.0

0852 8.0

1340 14.0

1340 14.0

1340 14.0

1340 14.0

TRAFFIC RATE HISTOGRAM

TIME OF DAY (HRS.)

ARRIVALS/HR.

DEPARTURES/HR.

TOTAL TRACKS/HR.

1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24
0
0
0 0

TRACKING DATA FOR: 6/28/80

ANALYSIS TYPE

ARRIVAL

DEPARTURE

TOTAL TRACKS

TARGET HISTOGRAM

NO. AIRCRAFT REPORTS 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20
FREQUENCY 8 29 55 69 53 24 29 13 6 1 0 0 0 0 0 0 0 0 0 0 0
PERCENTAGE 2 10 10 24 18 8 10 4 2 0 0 0 0 0 0 0 0 0 0 0 0

TOTAL TRACKS

RATE/HR.

START TIME

STOP TIME

29.

31.

68.

0847 8.0

0847 8.0

0847 8.0

1024 36.0

1024 36.0

1024 36.0

TRAFFIC RATE HISTOGRAM

TIME OF DAY (HRS.)

ARRIVALS/HR.

DEPARTURES/HR.

TOTAL TRACKS/HR.

1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24
0
0
0 0

TRACKING DATA FOR : 6/29/80

ANALYSIS TYPE		TOTAL TRACKS										RATE/HR.		START TIME		STOP TIME									
ARRIVAL		95										26.		1130 11.0		1515 24.3									
DEPARTURE		74										20.		1130 11.0		1515 24.3									
TOTAL TRACKS		210										57.		1130 11.0		1515 24.3									
TARGET HISTOGRAM																									
NO. AIRCRAFT REPORTS		0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20			
FREQUENCY		52	116	145	134	116	67	27	7	2	1	0	0	0	0	0	0	0	0	0	0	0			
PERCENTAGE		7	17	21	20	17	10	4	1	0	0	0	0	0	0	0	0	0	0	0	0	0			
TRAFFIC RATE HISTOGRAM																									
TIME OF DAY (HRS.)		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24
ARRIVALS/HR.		0	0	0	0	0	0	0	0	0	0	0	19	30	24	29	0	0	0	0	0	0	0	0	0
DEPARTURES/HR.		0	0	0	0	0	0	0	0	0	0	0	13	18	23	23	0	0	0	0	0	0	0	0	0
TOTAL TRACKS/HR.		0	0	0	0	0	0	0	0	0	0	0	43	48	55	72	0	0	0	0	0	0	0	0	0

TRACKING DATA FOR : 6/30/80

ANALYSIS TYPE		TOTAL TRACKS										RATE/HR.		START TIME		STOP TIME										
ARRIVAL		64										8.		0909 9.0		1703 31.8										
DEPARTURE		66										8.		0909 9.0		1703 31.8										
TOTAL TRACKS		198										25.		0909 9.0		1703 31.8										
TARGET HISTOGRAM																										
NO. AIRCRAFT REPORTS		0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24
FREQUENCY		458	509	288	116	36	16	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
PERCENTAGE		32	35	20	8	2	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
TRAFFIC RATE HISTOGRAM																										
TIME OF DAY (HRS.)		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	
ARRIVALS/HR.		0	0	0	0	0	0	0	0	0	7	7	17	3	4	7	14	6	0	0	0	0	0	0	0	0
DEPARTURES/HR.		0	0	0	0	0	0	0	0	0	8	8	16	6	2	8	12	7	0	0	0	0	0	0	0	0
TOTAL TRACKS/HR.		0	0	0	0	0	0	0	0	0	35	24	39	7	20	29	31	19	0	0	0	0	0	0	0	0

TRACKING DATA FOR: 7/1/80

ANALYSIS TYPE

ARRIVAL

DEPARTURE

TOTAL TRACKS

TARGET HISTOGRAM

TOTAL TRACKS

RATE/HR.

START TIME

STOP TIME

NO. AIRCRAFT REPORTS

FREQUENCY

PERCENTAGE

TRAFFIC RATE HISTOGRAM

TIME OF DAY (HRS.)

ARRIVALS/HR.

DEPARTURES/HR.

TOTAL TRACKS/HR.

TRACKING DATA FOR: 7/2/80

ANALYSIS TYPE

ARRIVAL

DEPARTURE

TOTAL TRACKS

TARGET HISTOGRAM

TOTAL TRACKS

RATE/HR.

START TIME

STOP TIME

NO. AIRCRAFT REPORTS

FREQUENCY

PERCENTAGE

TRAFFIC RATE HISTOGRAM

TIME OF DAY (HRS.)

ARRIVALS/HR.

DEPARTURES/HR.

TOTAL TRACKS/HR.

TRACKING DATA FOR: 7/3/80

ANALYSIS TYPE

ARRIVAL

DEPARTURE

TOTAL TRACKS

TARGET HISTOGRAM

NO. AIRCRAFT REPORTS	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
FREQUENCY	530	234	39	6	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
PERCENTAGE	65	28	4	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

TOTAL TRACKS

RATE/HR.

START TIME

STOP TIME

TRAFFIC RATE HISTOGRAM

TIME OF DAY (HRS.)

ARRIVALS/HR.

DEPARTURES/HR.

TOTAL TRACKS/HR.

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

TRACKING DATA FOR: 7/4/80

ANALYSIS TYPE

ARRIVAL

DEPARTURE

TOTAL TRACKS

TARGET HISTOGRAM

NO. AIRCRAFT REPORTS	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
FREQUENCY	38	141	169	226	122	47	27	7	5	2	0	0	0	0	0	0	0	0	0	0	0
PERCENTAGE	4	17	11	18	15	5	3	0	0	0	0	0	0	0	0	0	0	0	0	0	0

TOTAL TRACKS

RATE/HR.

START TIME

STOP TIME

TRAFFIC RATE HISTOGRAM

TIME OF DAY (HRS.)

ARRIVALS/HR.

DEPARTURES/HR.

TOTAL TRACKS/HR.

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

TRACKING DATA FOR: 7/5/80

ANALYSIS TYPE

ARRIVAL

DEPARTURE

TOTAL TRACKS

TARGET HISTOGRAM

NO. AIRCRAFT R

FREQUENCY

PERCENTAGE

TOTAL TRACKS

RATE/HR.

START TIME

STOP TIME

TRAFFIC RATE HISTOGRAM

TIME OF DAY (HRS.)

ARRIVALS/HR.

DEPARTURES/HR.

TOTAL TRACKS/HR.

TRACKING DATA FOR: 7/6/80

ANALYSIS TYPE

ARRIVAL

DEPARTURE

TOTAL TRACKS

TARGET HISTOGRAM

NO. AIRCRAFT R

FREQUENCY

PERCENTAGE

TOTAL TRACKS

RATE/HR.

START TIME

STOP TIME

TRAFFIC RATE HISTOGRAM

TIME OF DAY (HRS.)

ARRIVALS/HR.

DEPARTURES/HR.

TOTAL TRACKS/HR.

ANALYSIS TYPE	TOTAL TRACKS	RATE/HR.	START TIME	STOP TIME
ARRIVAL	107	13.	0855	1713
DEPARTURE	103	12.	0855	1713
TOTAL TRACKS	259	31.	0855	1713
TARGET HISTOGRAM				
NO. AIRCRAFT REPORTS	0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20			
FREQUENCY	301 568 381 152 37 30 14 11 0 0 0 0 0 0 0 0 0 0 0 0 0			
PERCENTAGE	20 38 25 10 2 2 0 0 0 0 0 0 0 0 0 0 0 0 0 0			
TRAFFIC RATE HISTOGRAM				
TIME OF DAY (HRS.)	1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24			
ARRIVALS/HR.	0 0 0 0 0 0 0 0 0 0 23 1 21 6 14 7 15 0 0 0 0 0 0			
DEPARTURES/HR.	0 0 0 0 0 0 0 0 0 0 14 6 17 7 13 12 11 16 0 0 0 0 0 0			
TOTAL TRACKS/HR.	0 0 0 0 0 0 0 0 0 0 34 17 48 23 29 30 28 34 0 0 0 0 0 0			

ANALYSIS TYPE	TOTAL TRACKS	RATE/HR.	START TIME	STOP TIME
ARRIVAL	177	22.	0840 8.0	1701 52.9
DEPARTURE	141	17.	0840 8.0	1701 52.9
TOTAL TRACKS	323	40.	0840 8.0	1701 52.9
TARGET HISTOGRAM				
NO. AIRCRAFT REPORTS	0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20			
FREQUENCY	184 433 405 265 114 50 10 3 0 0 0 0 0 0 0 0 0 0 0 0 0			
PERCENTAGE	12 29 27 18 7 3 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0			
TRAFFIC RATE HISTOGRAM				
TIME OF DAY (HRS.)	1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24			
ARRIVALS/HR.	0 0 0 0 0 0 0 0 0 0 22 19 22 16 19 15 41 18 0 0 0 0 0 0			
DEPARTURES/HR.	0 0 0 0 0 0 0 0 0 0 17 22 18 13 16 10 36 8 0 0 0 0 0 0			
TOTAL TRACKS/HR.	0 0 0 0 0 0 0 0 0 0 38 45 44 29 34 29 61 33 0 0 0 0 0 0			

TRACKING DATA FOR : 7/10/80

ANALYSIS TYPE

ARRIVAL

DEPARTURE

TOTAL TRACKS

TARGET HISTOGRAM

NO. AIRCRAFT REPORTS

FREQUENCY

PERCENTAGE

TOTAL TRACKS

125

105

283

RATE/HR.

15.

13.

35.

START TIME

0852 8.0

0852 8.0

0852 8.0

STOP TIME

1700 22.3

1700 22.3

1700 22.3

TRAFFIC RATE HISTOGRAM

TIME OF DAY (HRS.)

ARRIVALS/HR.

DEPARTURES/HR.

TOTAL TRACKS/HR.

1 2 3 4 5 6 7 8 9 10

0 0 0 0 0 0 0 0 0 0

0 0 0 0 0 0 0 0 0 0

0 0 0 0 0 0 0 0 0 0

0 0 0 0 0 0 0 0 0 0

11 12 13 14 15 16 17 18 19 20

0 0 0 0 0 0 0 0 0 0

0 0 0 0 0 0 0 0 0 0

0 0 0 0 0 0 0 0 0 0

0 0 0 0 0 0 0 0 0 0

21 22 23 24

0 0 0 0

0 0 0 0

0 0 0 0

0 0 0 0

TRACKING DATA FOR : 7/11/80

ANALYSIS TYPE

ARRIVAL

DEPARTURE

TOTAL TRACKS

TARGET HISTOGRAM

NO. AIRCRAFT REPORTS

FREQUENCY

PERCENTAGE

TOTAL TRACKS

108

100

262

RATE/HR.

13.

12.

31.

START TIME

0848 8.0

0848 8.0

0848 8.0

STOP TIME

1721 44.7

1721 44.7

1721 44.7

TRAFFIC RATE HISTOGRAM

TIME OF DAY (HRS.)

ARRIVALS/HR.

DEPARTURES/HR.

TOTAL TRACKS/HR.

1 2 3 4 5 6 7 8 9 10

0 0 0 0 0 0 0 0 0 0

0 0 0 0 0 0 0 0 0 0

0 0 0 0 0 0 0 0 0 0

0 0 0 0 0 0 0 0 0 0

11 12 13 14 15 16 17 18 19 20

0 0 0 0 0 0 0 0 0 0

0 0 0 0 0 0 0 0 0 0

0 0 0 0 0 0 0 0 0 0

0 0 0 0 0 0 0 0 0 0

21 22 23 24

0 0 0 0

0 0 0 0

0 0 0 0

0 0 0 0

TRACKING DATA FOR: 7/14/80

ANALYSIS TYPE

ARRIVAL

DEPARTURE

TOTAL TRACKS

TARGET HISTOGRAM

NO. AIRCRAFT REPORTS 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20
FREQUENCY 216 439 318 311 136 46 12 4 2 0 0 0 0 0 0 0 0 0 0 0 0
PERCENTAGE 14 29 21 20 9 3 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0

TOTAL TRACKS

RATE/HR.

START TIME

STOP TIME

TRAFFIC RATE HISTOGRAM

TIME OF DAY (HRS.)

ARRIVALS/HR.

DEPARTURES/HR.

TOTAL TRACKS/HR.

1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24
0
0
0 0

TRACKING DATA FOR: 7/15/80

ANALYSIS TYPE

ARRIVAL

DEPARTURE

TOTAL TRACKS

TARGET HISTOGRAM

NO. AIRCRAFT REPORTS 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20
FREQUENCY 129 366 418 273 155 64 37 5 2 0 0 0 0 0 0 0 0 0 0 0 0
PERCENTAGE 8 25 28 18 10 4 2 0 0 0 0 0 0 0 0 0 0 0 0 0 0

TOTAL TRACKS

RATE/HR.

START TIME

STOP TIME

TRAFFIC RATE HISTOGRAM

TIME OF DAY (HRS.)

ARRIVALS/HR.

DEPARTURES/HR.

TOTAL TRACKS/HR.

1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24
0
0
0 0

TRACKING DATA FOR: 7/16/80

ANALYSIS TYPE

ARRIVAL

DEPARTURE

TOTAL TRACKS

TARGET HISTOGRAM

TOTAL TRACKS

RATE/HR.

START TIME

STOP TIME

NO. AIRCRAFT REPORTS	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
FREQUENCY	307	555	319	119	61	10	2	1	0	0	0	0	0	0	0	0	0	0	0	0	0
PERCENTAGE	22	40	23	8	4	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

TRAFFIC RATE HISTOGRAM

TIME OF DAY (HRS.)

ARRIVALS/HR.

DEPARTURES/HR.

TOTAL TRACKS/HR.

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24
0	0	0	0	0	0	0	0	0	0	0	15	14	0	14	6	6	13	22	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	13	11	8	0	7	7	4	10	6	0	0	0	0	0
0	0	0	0	0	0	0	0	0	25	35	35	35	0	35	19	18	24	42	0	0	0	0	0

TRACKING DATA FOR: 7/17/80

ANALYSIS TYPE

ARRIVAL

DEPARTURE

TOTAL TRACKS

TARGET HISTOGRAM

TOTAL TRACKS

RATE/HR.

START TIME

STOP TIME

NO. AIRCRAFT REPORTS	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
FREQUENCY	496	394	242	102	19	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
PERCENTAGE	39	31	19	8	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

TRAFFIC RATE HISTOGRAM

TIME OF DAY (HRS.)

ARRIVALS/HR.

DEPARTURES/HR.

TOTAL TRACKS/HR.

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24
0	0	0	0	0	0	0	0	0	13	10	4	8	13	5	12	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	16	6	8	3	8	4	8	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	27	36	22	22	12	28	12	22	0	0	0	0	0	0	0	0

TRACKING DATA FOR: 7/18/80

ANALYSIS TYPE

ARRIVAL

DEPARTURE

TOTAL TRACKS

TARGET HISTOGRAM

NO. AIRCRAFT REPORTS 0

FREQUENCY

PERCENTAGE

TOTAL TRACKS

RATE/HR.

START TIME

STOP TIME

TRAFFIC RATE HISTOGRAM

TIME OF DAY (HRS.)

ARRIVALS/HR.

DEPARTURES/HR.

TOTAL TRACKS/HR.

11 12 13

14 15 16

17 18 19

TRACKING DATA FOR: 7/19/80

ANALYSIS TYPE

ARRIVAL

DEPARTURE

TOTAL TRACKS

TARGET HISTOGRAM

NO. AIRCRAFT REPORTS 0

FREQUENCY

PERCENTAGE

TOTAL TRACKS

RATE/HR.

START TIME

STOP TIME

TRAFFIC RATE HISTOGRAM

TIME OF DAY (HRS.)

ARRIVALS/HR.

ARRIVALS/HR.
DEPARTURES/HR.

TOTAL TRACKS/HR.

11 12 13

14 15 16

17 18 19

TRACKING DATA FOR: 7/20/80

ANALYSIS TYPE

ARRIVAL

DEPARTURE

TOTAL TRACKS

TARGET HISTOGRAM

NO. AIRCRAFT REPORTS

NO. AIRCRAFT
FREQUENCY

PERCENTAGE

TOTAL TRACKS

RATE/HR.

START TIME

STOP TIME

TRAFFIC RATE HISTOGRAM

TIME OF DAY (HRS.)

ARRIVALS/HR.

DEPARTURES/HR.

TOTAL TRACKS/HR.

TRACKING DATA FOR: 7/21/80

ANALYSIS TYPE

ARRIVAL

DEPARTURE

TOTAL TRACKS

TARGET HISTOGRAM

NO. AIRCRAFT REPORTS

FREQUENCY

PERCENTAGE

TOTAL TRACKS

RATE/HR.

START TIME

STOP TIME

TRAFFIC RATE HISTOGRAM

TIME OF DAY (HRS.)

TIME OF DAY (ARRIVALS/HR.

DEPARTURES/HR.

TOTAL TRACKS/HR.

TRACKING DATA FOR: 7/22/80

ANALYSIS TYPE

ARRIVAL

DEPARTURE

TOTAL TRACKS

TARGET HISTOGRAM

TOTAL TRACKS

RATE/HR.

START TIME

STOP TIME

NO. AIRCRAFT REPORTS

FREQUENCY

PERCENTAGE

TRAFFIC RATE HISTOGRAM

TIME OF DAY (HRS.)

ARRIVALS/HR.

DEPARTURES/HR.

TOTAL TRACKS/HR.

TRACKING DATA FOR: 7/24/80

ANALYSIS TYPE

ARRIVAL

DEPARTURE

TOTAL TRACKS

TARGET HISTOGRAM

TOTAL TRACKS

RATE/HR.

START TIME

STOP TIME

NO. AIRCRAFT REPORTS

FREQUENCY

PERCENTAGE

TRAFFIC RATE HISTOGRAM

TIME OF DAY (HRS.)

ARRIVALS/HR.

DEPARTURES/HR.

TOTAL TRACKS/HR.

TIME OF DAY (HRS.)	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	
ARRIVALS/HR.	0	0	0	0	0	0	0	0	0	0	7	25	14	11	6	0	0	4	0	0	0	0	0	0	0
DEPARTURES/HR.	0	0	0	0	0	0	0	0	0	10	22	10	13	5	0	0	0	0	0	0	0	0	0	0	0
TOTAL TRACKS/HR.	0	0	0	0	0	0	0	0	0	20	56	49	30	22	0	0	0	5	0	0	0	0	0	0	0

TIME OF DAY (HRS.)	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24
ARRIVALS/HR.	0	0	0	0	0	0	0	0	0	0	8	14	8	0	0	0	0	0	0	0	0	0	0	0
DEPARTURES/HR.	0	0	0	0	0	0	0	0	0	9	13	15	0	0	0	0	0	0	0	0	0	0	0	0
TOTAL TRACKS/HR.	0	0	0	0	0	0	0	0	0	27	36	28	0	0	0	0	0	0	0	0	0	0	0	0

TRACKING DATA FOR : 7/28/80

ANALYSIS TYPE

ARRIVAL

DEPARTURE

TOTAL TRACKS

TARGET HISTOGRAM

NO. AIRCRAFT REPORTS

FREQUENCY

PERCENTAGE

TOTAL TRACKS

RATE/HR.

START TIME

STOP TIME

TRAFFIC RATE HISTOGRAM

TIME OF DAY (HRS.)

ARRIVALS/HR.

DEPARTURES/HR.

TOTAL TRACKS/HR.

TRACKING DATA FOR : 7/29/80

ANALYSIS TYPE

ARRIVAL

DEPARTURE

TOTAL TRACKS

TARGET HISTOGRAM

NO. AIRCRAFT REPORTS

FREQUENCY

PERCENTAGE

TOTAL TRACKS

RATE/HR.

START TIME

STOP TIME

TRAFFIC RATE HISTOGRAM

TIME OF DAY (HRS.)

ARRIVALS/HR.

DEPARTURES/HR.

TOTAL TRACKS/HR.

TRACKING DATA FOR: 7/30/80

ANALYSIS TYPE

ARRIVAL

DEPARTURE

TOTAL TRACKS

TARGET HISTOGRAM

NO. AIRCRAFT REPORTS 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20
 FREQUENCY 242 288 331 256 129 62 18 3 1 0 0 0 0 0 0 0 0 0 0 0 0
 PERCENTAGE 18 21 24 19 9 4 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0

TOTAL TRACKS

RATE/HR.

START TIME

STOP TIME

150
121
320
20.
16.
43.
0911 9.0
0911 9.0
0911 9.0

TRAFFIC RATE HISTOGRAM

TIME OF DAY (HRS.)

ARRIVALS/HR.

DEPARTURES/HR.

TOTAL TRACKS/HR.

1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24
 0 0 0 0 0 0 0 0 0 0 15 12 28 30 23 22 20 11 0 0 0 0 0
 0 0 0 0 0 0 0 0 0 0 11 13 23 20 16 26 6 12 0 0 0 0 0
 0 0 0 0 0 0 0 0 0 0 35 33 49 57 48 56 29 21 0 0 0 0 0

TRACKING DATA FOR: 8/1/80

ANALYSIS TYPE

ARRIVAL

DEPARTURE

TOTAL TRACKS

TARGET HISTOGRAM

NO. AIRCRAFT REPORTS 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20
 FREQUENCY 385 334 244 113 31 3 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
 PERCENTAGE 34 30 21 10 2 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0

TOTAL TRACKS

RATE/HR.

START TIME

STOP TIME

56
55
166
9.
9.
27.
0911 9.0
0911 9.0
0911 9.0

TRAFFIC RATE HISTOGRAM

TIME OF DAY (HRS.)

ARRIVALS/HR.

DEPARTURES/HR.

TOTAL TRACKS/HR.

1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24
 0 0 0 0 0 0 0 0 0 4 8 15 12 2 0 1 13 0 0 0 0 0 0
 0 0 0 0 0 0 0 0 0 3 8 18 8 5 0 3 13 0 0 0 0 0 0
 0 0 0 0 0 0 0 0 0 24 31 43 31 18 0 8 22 0 0 0 0 0 0

TRACKING DATA FOR: 8/3/80

ANALYSIS TYPE

ARRIVAL

DEPARTURE

TOTAL TRACKS

TARGET HISTOGRAM

NO. AIRCRAFT REPORTS

FREQUENCY

PERCENTAGE

TOTAL TRACKS

RATE/HR.

START TIME

STOP TIME

TRAFFIC RATE HISTOGRAM

TIME OF DAY (HRS.)

ARRIVALS/HR.

DEPARTURES/HR.

TOTAL TRACKS/HR.

TRACKING DATA FOR: 8/11/80

ANALYSIS TYPE

ARRIVAL

DEPARTURE

TOTAL TRACKS

TARGET HISTOGRAM

NO. AIRCRAFT REPORTS

FREQUENCY

PERCENTAGE

TOTAL TRACKS

RATE/HR.

START TIME

STOP TIME

TRAFFIC RATE HISTOGRAM

TIME OF DAY (HRS.)

ARRIVALS/HR.

DEPARTURES/HR.

TOTAL TRACKS/HR.

TRACKING DATA FOR: 8/12/80

ANALYSIS TYPE

ARRIVAL

DEPARTURE

TOTAL TRACKS

TARGET HISTOGRAM

NO. AIRCRAFT R

FREQUENCY

PERCENTAGE

TOTAL TRACKS

RATE/HR.

START TIME

STOP TIME

TRAFFIC RATE HISTOGRAM

TIME OF DAY (HRS.)

ARRIVALS/HR.

DEPARTURES/HR.

TOTAL TRACKS/HR.

17 12 12

10

1

TRACKING DATA FOR: 8/13/80

ANALYSIS TYPE

ARRIVAL

DEPARTURE

TOTAL TRACKS

TARGET HISTOGRAM

NO. AIRCRAFT R

NO. HINCKLEY N
FREQUENCY

PERCENTAGE

TOTAL TRACKS

RATE/HR.

START TIME

STOP TIME

TRAFFIC RATE HISTOGRAM

TIME OF DAY (HRS.)

ARRIVAL S/HR.

ARRIVALS/HR.
DEPARTURES/HR.

TOTAL TRACKS/HR.

11 12 13

14 1E 1

2

TRACKING DATA FOR: 8/14/80

ANALYSIS TYPE

ARRIVAL

DEPARTURE

TOTAL TRACKS

TARGET HISTOGRAM

NO. AIRCRAFT REPORTS

FREQUENCY

PERCENTAGE

TOTAL TRACKS

RATE/HR.

START TIME

STOP TIME

TRAFFIC RATE HISTOGRAM

TIME OF DAY (HRS.)

ARRIVALS/HR.

DEPARTURES/HR.

TOTAL TRACKS/HR.

TRACKING DATA FOR: 8/15/80

ANALYSIS TYPE

ARRIVAL

DEPARTURE

TOTAL TRACKS

TARGET HISTOGRAM

NO. AIRCRAFT REPORTS

FREQUENCY

PERCENTAGE

TOTAL TRACKS

RATE/HR.

START TIME

STOP TIME

TRAFFIC RATE HISTOGRAM

TIME OF DAY (HRS.):

ARRIVALS/HR.

DEPARTURES/HR.

TOTAL TRACKS/HR.

TRACKING DATA FOR: 8/16/80

ANALYSIS TYPE

ARRIVAL

DEPARTURE

TOTAL TRACKS

TARGET HISTOGRAM

NO. AIRCRAFT REPORTS 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20
 FREQUENCY 379 578 530 365 214 80 20 6 3 0 1 0 0 0 0 0 0 0 0 0 0
 PERCENTAGE 17 26 24 16 9 3 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0

TOTAL TRACKS

203

197

469

RATE/HR.

17.

16.

39.

START TIME

0905 9.0

0905 9.0

0905 9.0

STOP TIME

2158 9.6

2158 9.6

2158 9.6

TRAFFIC RATE HISTOGRAM

TIME OF DAY (HRS.)

ARRIVALS/HR.

DEPARTURES/HR.

TOTAL TRACKS/HR.

1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24
 0
 0
 0

REFERENCES

1. Parker, L. C., et al.: "An Automated Pilot Advisory System," (A paper presented at the 1978 IEEE Region 3 Conference and Exhibit in Atlanta, GA, April 10-12, 1978).
2. Haidt, J. G., et al.: "Pilot Advisory System Feasibility Definition Study," Research Triangle Institute, February 1977.
3. Parks, John L., Jr., "Automated Pilot Advisory System Test and Evaluation at Manassas Municipal Airport," 1980.

1. Report No. NASA TM-73296		2. Government Accession No.		3. Recipient's Catalog No.	
4. Title and Subtitle AUTOMATED PILOT ADVISORY SYSTEM				5. Report Date October 1981	
				6. Performing Organization Code 024.1	
7. Author(s) John L. Parks, Jr. NASA Wallops Flight Center Wallops Island, VA 23337				8. Performing Organization Report No.	
Mr. James G. Haidt Research Triangle Institute Research Triangle Park, NC 27709				10. Work Unit No.	
9. Performing Organization Name and Address National Aeronautics and Space Administration Wallops Flight Center Wallops Island, VA 23337				11. Contract or Grant No.	
				13. Type of Report and Period Covered Technical Memorandum	
12. Sponsoring Agency Name and Address National Aeronautics and Space Administration Washington, DC 20546				14. Sponsoring Agency Code	
15. Supplementary Notes Final Report					
16. Abstract An Automated Pilot Advisory System (APAS) was developed and operationally tested to demonstrate the concept that low cost automated systems could provide air traffic and aviation weather advisory information at high density uncontrolled airports. The system was designed to enhance the see-and-be-seen rule of flight, and pilots who used the system preferred it over the self-announcement system presently used at uncontrolled airports.					
17. Key Words (Suggested by Author(s)) Air traffic advisories Track while scan Aviation weather advisory Radar Low cost General aviation Automated system Uncontrolled airports				18. Distribution Statement Unclassified - unlimited STAR Category 04	
19. Security Classif. (of this report) Unclassified		20. Security Classif. (of this page) Unclassified		21. No. of Pages 179	
				22. Price*	

National Aeronautics and
Space Administration

SPECIAL FOURTH CLASS MAIL
BOOK

Postage and Fees Paid
National Aeronautics and
Space Administration
NASA-451



Wallops Island, VA
23337

Official Business

Penalty for Private Use, \$300

NASA

POSTMASTER: If Undeliverable (Section 158
Postal Manual) Do Not Return
